EXHIBIT C

2022-1866, 2022-1867, 2022-1868

United States Court of Appeals

for the

Federal Circuit

GOOGLE LLC,
Appellant,

v.

SINGULAR COMPUTING LLC,
Appellee.

Appeals from the United States Patent and Trademark Office,
Patent Trial and Appeal Board in
Case Nos. IPR2021-00155, IPR2021-00165, and IRP2021-00179

CORRECTED NON-CONFIDENTIAL RESPONSE BRIEF OF APPELLEE, SINGULAR COMPUTING LLC

Alfred R. Fabricant
Peter Lambrianakos
Vincent J. Rubino, III
Richard M. Cowell
FABRICANT LLP
411 Theodore Fremd Avenue
Suite 206 South
Rye, New York 10580
(212) 257-5797

James J. Foster
PRINCE LOBEL TYE LLP
One International Place
Suite 3700
Boston, MA 02110
(617) 456-8000

December 23, 2022

Counsel for Appellee Singular Computing LLC

PATENT CLAIMS AT ISSUE

U.S. Patent No. 10,416,961

1. A device comprising:

at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,

wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/65,000 through 65,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X % of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input; and at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit.

2. The device of claim 1, wherein the at least one first computing device comprises at least one of a central processing unit (CPU), a graphics processing unit (GPU), a field programmable gate array (FPGA), a microcode-based processor, a hardware sequencer, and a state machine.

3. The device of claim 2, wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

21. A device comprising a computer processor and a computer-readable memory storing computer program instructions, wherein the computer program instructions are executable by the processor to emulate a second device comprising:

a plurality of components comprising:

at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing

a first numerical value to produce a first output signal representing a second numerical value;

wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/65,000 through 65,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X % of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input.

23. The device of claim 21, wherein the number of LPHDR execution units in the second device exceeds by at least one hundred the non-negative integer number of execution units in the second device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

PATENT CLAIMS AT ISSUE

U.S. Patent No. 9,218,156

1. A device comprising:

at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,

wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/65,000 through 65,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input;

and at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit.

2. The method of claim 1, wherein the at least one first computing device comprises at least one of a central processing unit (CPU), a graphics processing unit (GPU), a field programmable gate array (FPGA), a microcode-based processor, a hardware sequencer, and a state machine.

3. The device of claim 2, wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

4. The device of claim 3, wherein X=10%.

5. The device of claim 3, wherein Y=0.2%.

6. The device of claim 3, wherein X=10% and Y=0.2%.

7. The device of claim 3, wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000.

8. The device of claim 3, wherein the first operation is multiplication.

PATENT CLAIMS AT ISSUE

U.S. Patent No. 8,407,273

1. A device:

comprising at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,

wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/65,000 through 65,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X % of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input.

3. The device of claim 1, wherein the at least one first LPHDR execution unit comprises at least ten LPHDR execution units.

4. The device of claim 3, wherein the at least one first LPHDR execution unit comprises at least part of an FPGA.

5. The device of claim 1, wherein the number of LPHDR execution units in the device exceeds by at least ten the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

6. The device of claim 5, wherein the at least one first LPHDR execution unit comprises at least part of an FPGA.

7. The device of claim 1, wherein the at least one first LPHDR execution unit comprises at least one hundred LPHDR execution units.

8. The device of claim 1, wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

9. The device of claim 1, wherein the at least one first LPHDR execution unit comprises at least five hundred LPHDR execution units.

10. The device of claim 1, wherein the number of LPHDR execution units in the device exceeds by at least five hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

11. The device of claim 8, wherein X=10%.

12. The device of claim 8, wherein Y=0.1%.

13. The device of claim 8, wherein Y=0.15%.

14. The device of claim 8, wherein Y=0.2%.

15. The device of claim 8, wherein X=10% and wherein Y=0.1%.

16. The device of claim 8, wherein X=10% and wherein Y=0.15%.

17. The device of claim 8, wherein X=10% and wherein Y=0.2%.

18. The device of claim 8, wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000.

19. The device of claim 1, wherein the at least one first LPHDR execution unit comprises a plurality of locally connected LPHDR execution units.

20. The device of claim 1, wherein the device has a SIMD architecture.

25. The device of claim 1, wherein the at least one LPHDR execution unit comprises at least five hundred locally connected LPHDR execution units, wherein the device includes memory

33. A device comprising a computer processor and a computer-readable memory storing computer program instructions, wherein the computer program instructions are executable by the processor to emulate a second device comprising:

at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value;

wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/65,000 through 65,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X % of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input.

34. The device of claim 33, wherein the at least one first LPHDR execution unit comprises at least ten LPHDR execution units.

35. The device of claim 33, wherein the number of LPHDR execution units in the second device exceeds by at least ten the non-negative integer number of execution units in the second device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

36. A device:

comprising at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,

wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/65,000 through 65,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X % of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input;

wherein the number of LPHDR execution units in the device exceeds the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

37. The device of claim 36, wherein the at least one first LPHDR execution unit comprises at least part of an FPGA.

38. The device of claim 36, wherein the at least one first LPHDR execution unit comprises at least ten LPHDR execution units.

39. The device of claim 38, wherein the at least one first LPHDR execution unit comprises at least part of an FPGA.

40. The device of claim 36, wherein the number of LPHDR execution units in the device exceeds by at least ten the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

41. The device of claim 40, wherein the at least one first LPHDR execution unit comprises at least part of an FPGA.

42. The device of claim 36, wherein the at least one first LPHDR execution unit comprises at least one hundred LPHDR execution units.

43. The device of claim 36, wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

44. The device of claim 36, wherein the at least one first LPHDR execution unit comprises at least five hundred LPHDR execution units.

45. The device of claim 36, wherein the number of LPHDR execution units in the device exceeds by at least five hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

46. The device of claim 43, wherein X=10%.

47. The device of claim 43, wherein Y=0.1%.

48. The device of claim 43, wherein Y=0.15%.

49. The device of claim 43, wherein Y=0.2%.

50. The device of claim 43, wherein X=10% and wherein Y=0.1%. 51.

51. The device of claim 43, wherein X=10% and wherein Y=0.15%.

52. The device of claim 43, wherein X=10% and wherein Y=0.2%.

53. The device of claim 43, wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000.

54. The device of claim 36, wherein the at least one first LPHDR execution unit comprises a plurality of locally connected LPHDR execution units.

55. The device of claim 36, wherein the device has a SIMD architecture.

56. The device of claim 36, wherein the device includes memory locally accessible to the at least one first LPHDR execution unit.

57. The device of claim 36, wherein the device is implemented on a silicon chip.

58. The device of claim 36, wherein the device is implemented on a silicon chip using digital technology.

59. The device of claim 36, wherein the device further comprises a digital processor adapted to control the operation of the at least one first LPHDR execution unit.

60. The device of claim 36, wherein the at least one LPHDR execution unit comprises at least five hundred locally connected LPHDR execution units, wherein the device includes memory locally accessible to at least one of the LPHDR execution units, and wherein the device is implemented on a silicon chip using digital technology.

61. The device of claim 36, wherein the device is part of a mobile device.

63. The device of claim 36, wherein the at least one first LPHDR execution unit represents numbers using a floating point representation.

67. The device of claim 36, wherein the device is adapted to perform nearest neighbor search.

68. A device comprising a computer processor and a computer-readable memory storing computer program instructions, wherein the computer program instructions are executable by the processor to emulate a second device comprising:

least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,

wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/65,000 through 65,000 and for at least 5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least 5% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least 0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input;

wherein the number of LPHDR execution units in the second device exceeds the non-negative integer number of execution units in the second device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

69. The device of claim 68, wherein the at least one first LPHDR execution unit comprises at least ten LPHDR execution units.

70. The device of claim 68, wherein the number of LPHDR execution units in the second device exceeds by at least ten the non-negative integer number of execution units in the second device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

CERTIFICATE OF INTEREST

Counsel for Singular Computing LLC certifies the following:

1. Provide the full names of all entities represented by undersigned counsel in this case:

Singular Computing LLC

2. Provide the full names of all real parties in interest (if the party named in the caption is not the real party in interest) for the entities. Do not list the real parties if they are the same as the entities:

None/Not Applicable

3. Provide the full names of all parent corporations for the entities and all publicly held companies that own 10 percent or more of the stock in the entities:

None/Not Applicable

4. The names of all law firms and the partners or associates that (a) appeared for the entities in the originating court or agency or (b) are expected to appear in this court for the entities. Do not include those who have already entered an appearance in this court. Fed. Cir. R. 47.4(a)(4):

FABRICANT LLP, 411 Theodore Fremd Avenue, Suite 206 South, Rye, New York 10580: Enrique Iturralde

PRINCE LOBEL TYE LLP, 1 International Place, Suite 37000, Boston, Massachusetts 02110: Matthew Vella, Brian Seeve

5. The following cases are pending in a court or agency and will directly affect or be directly affected by this court's decision in the pending appeal:

Singular Computing LLC v. *Google* LLC, No. 1:19-cv-12551 (D. Mass.)

6. Provide any information required, Case No under Fed. R. App. P. 26.1(b) (organizational victims in criminal cases) and 26.1(c) (bankruptcy case debtors and trustees). Fed. Cir. R. 47.4(a)(6).

None/Not Applicable

Dated: December 23, 2022

By: */s/ Alfred R. Fabricant*
Alfred R. Fabricant
Fabricant LLP

TABLE OF CONTENTS

Page(s)

CERTIFICATE OF INTEREST xi

TABLE OF AUTHORITIES xvii

STATEMENT OF RELATED CASES xx

TABLE OF ABBREVIATIONS xxi

STATEMENT OF THE ISSUES 1

INTRODUCTION 2

STATEMENT OF THE CASE 5

I. THE CHALLENGED PATENTS 5

A. The '961 Patent 5

B. The '156 Patent 9

C. The '273 Patent 9

II. OVERVIEW OF THE CITED REFERENCES 10

A. Dockser 10

B. Tong 13

C. MacMillan 15

III. THE IPR PROCEEDINGS 17

A. The Board's Determination that a POSA Would Not Combine Dockser and MacMillan 18

B. The Board's Determination that a POSA Would Not Combine Dockser, Tong, and MacMillan 19

IV. PROCEEDINGS IN THIS COURT 20

SUMMARY OF THE ARGUMENT 22

STANDARD OF REVIEW 23

ARGUMENT 25

V. THE BOARD'S FINDING THAT A POSA WOULD NOT HAVE BEEN MOTIVATED TO COMBINE DOCKSER WITH MACMILLAN BECAUSE IT WOULD HAVE DESTROYED MACMILLAN'S STATED OBJECTIVE OF AVOIDING INCREASED COST WAS SUPPORTED BY SUBSTANTIAL EVIDENCE.....25

A. Dockser and MacMillan Were Directed to Different Problems.....26

B. The Board's Finding Was Based on Substantial Evidence as to These Differences. 27

C. The Board Properly Considered that Combining Dockser with MacMillan Would Increase Cost Because of Added Complexity.....28

D. Substantial Evidence Supports the Board's Finding that Adding Dockser to MacMillan Would Increase Macmillan's Complexity, Size, and Therefore Cost 32

E. The Board Did Not Commit a Legal Error by Finding MacMillan Teaches Away from Including Dockser 36

VI. THE BOARD'S FINDING THAT A POSA WOULD NOT HAVE BEEN MOTIVATED TO COMBINE A CUSTOMIZED VERSION OF DOCKSER IN THE DOCKSER-MACMILLAN-TONG COMBINATION WAS SUPPORTED BY SUBSTANTIAL EVIDENCE.....40

A. Dockser Teaches a Processor that Operates at Traditional Precision with the Ability to Select a Lower Precision Based on the Application to Be Executed 40

B. Tong Does Not Support the Dockser-Tong-MacMillan Combination Because It Teaches that Full Precision Is Necessary 45

C. A POSA Would Not Modify Dockser as Google Suggests 47

D. Google Relied on Economic Efficiencies as the Motivation to Combine While Simultaneously Arguing that Economic Concerns Are Irrelevant as a Matter of Law....51

VII. THE BOARD'S DECISION IS FURTHER SUPPORTED BY STRONG EVIDENCE OF OBJECTIVE INDICIA OF NON-OBVIOUSNESS 52

A. Google's Initial Skepticism of the Invention Shows Non-Obviousness 54

B. Praise from Leaders in the Field Shows Non-Obviousness 56

C. Google's Copying of the Invention Shows Non-Obviousness 57

VIII. ALTERNATIVELY, THE COURT SHOULD AFFIRM THE BOARD'S DECISION AS TO GROUNDS 3 AND 4 BECAUSE THE COMBINATION DOES NOT DISCLOSE OR RENDER OBVIOUS THE "EXCEEDS" LIMITATION 60

A. The Dockser-MacMillan Combination Does Not Disclose Systems Where the Number of LPHDR Execution Units Exceeds the Number of Execution Units Adapted to Perform Full-Precision Multiplication....62

1. Dockser's FPP Is a Traditional-Precision Execution Unit Because It Is Adapted to Execute at Least the Operation of Multiplication on Floating Point Numbers that Are at Least 32 Bits Wide....63

2. Google's Interpretation of "Adapted to Execute at Least the Operation of Multiplication on Floating Point Numbers that Are at Least 32 bits Wide" Is Inconsistent with the Ordinary Meaning of the Exceeds Limitation and the Specification 65

CONCLUSION 67

CERTIFICATE OF COMPLIANCE 68

TABLE OF AUTHORITIES

Page(s)

Cases

Amazon.com, Inc. v. Barnesandnoble.com, Inc.,
239 F.3d 1343 (Fed. Cir. 2001) 23

Apple Inc. v. Samsung Elecs. Co.,
839 F.3d 1034 (Fed. Cir. 2016) (*en banc*) 53

Applied Med. Res. Corp. v. U.S. Surgical Corp.,
448 F.3d 1324 (Fed. Cir. 2006) 65

AstraZeneca AB v. Mylan Pharms. Inc.,
19 F.4th 1325 (Fed. Cir. 2021) 36

Chemours Co. FC, LLC v. Daikin Indus., Ltd.,
4 F.4th 1370 (Fed. Cir. 2021) 27, 37

Consol. Edison Co. v. NLRB,
305 U.S. 197 (1938).....23

Corephotonics, Ltd. v. Apple Inc.,
No. 2020-1961, 2021 WL 4944471 (Fed. Cir. 2021) 38

In re Dembiczak,
175 F.3d 994 (Fed. Cir. 1999) 23

In re Farrenkopf,
713 F.2d 714 (Fed. Cir. 1983) 29, 30

Friskit, Inc. v. Real Networks, Inc.,
306 F. App'x 610 (Fed. Cir. 2009).....31

In re Gartside,
203 F.3d 1305 (Fed. Cir. 2000) 23

Google LLC v. IPA Tech. Inc.,
34 F.4th 1081 (Fed. Cir. 2022) 24, 32

In re Gordon,
733 F.2d 900 (Fed. Cir. 1984) 37

Grit Energy Sols., LLC v. Oren Techs., LLC,
957 F.3d 1309 (Fed. Cir. 2020) 31, 32

Henny Penny Corp. v. Frymaster LLC,
938 F.3d 1324 (Fed. Cir. 2019) 38

Intel Corp. v. Qualcomm Inc.,
21 F.4th 784 (Fed. Cir. 2021) 38

K/S Himpp v. Hear-Wear Techs., LLC,
751 F.3d 1362 (Fed. Cir. 2014) 25

In re Kahn,
441 F.3d 977 (Fed. Cir. 2006) 39, 42

Liqwd, Inc. v. L'Oreal USA, Inc.,
941 F.3d 1133 (Fed. Cir. 2019) 60

In re Magna Elecs., Inc.,
611 F. App'x 969 (Fed. Cir. 2015) 31

In re Man Mach. Interface Techs. LLC,
822 F.3d 1282 (Fed. Cir. 2016) 64

Mintz v. Dietz & Watson, Inc.,
679 F.3d 1372 (Fed. Cir. 2012) 56

In re Mouttet,
686 F.3d 1322 (Fed. Cir. 2012) 39

Mylan Pharms. Inc. v. Merck Sharp & Dohme Corp.,
50 F.4th 147 (Fed. Cir. 2022) 23

Neptune Generics, LLC v. Eli Lilly & Co.,
921 F.3d 1372 (Fed. Cir. 2019) 54

In re Nilssen,
837 F.2d 1098 (Fed. Cir. 1987) 31

In re Nuvasive, Inc.,
842 F.3d 1376 (Fed. Cir. 2016) 44, 45, 46

Orthopedic Equip. Co., Inc. v. United States,
702 F.2d 1005 (Fed. Cir. 1983) 31

Outdry Techs. Corp. v. Geox S.p.A.,
859 F.3d 1364 (Fed. Cir. 2017) 25

Polaris Innovations, Ltd. v. Brent,
48 F.4th 1365 (Fed. Cir. 2022) 23

Power Integrations, Inc. v. Fairchild Semiconductor Int'l, Inc.
711 F.3d 1348 (Fed. Cir. 2013) 60

Provisur Techs., Inc. v. Weber, Inc.,
No. 21-1851, 2022 WL 17688071 (Fed. Dec. 15, 2022).....35

Redline Detection, LLC v. Star Envirotech, Inc.,
811 F.3d 435 (Fed. Cir. 2015) 25, 47

Trivascular, Inc. v. Samuels,
812 F.3d 1056 (Fed. Cir. 2016) 27

In re Watts,
354 F.3d 1362 (Fed. Cir. 2004) 45

WBIP, LLC v. Kohler Co.,
829 F.3d 1317 (Fed. Cir. 2016) 54, 56, 57

CONFIDENTIAL MATERIAL OMITTED

The Material redacted on page 56 of this Brief is subject to the Board's Protective Orders issued in IPR2021-00155, IPR2021-00165 and IPR2021-00179. The Material contains quotations from a confidential email subject to a non-disclosure agreement and the identity of the email's author.

STATEMENT OF RELATED CASES

In accordance with Federal Circuit Rule 47.5, the following case involving U.S. Patent Nos. 8,407,273 (the "'273 Patent"), 9,218,156 (the "'156 Patent"), and 10,416,961 (the "'961 Patent") may directly affect or be directly affected by a decision in these appeals: *Singular Computing LLC v. Google LLC*, 19-cv-12551-FDS (D. Mass).

TABLE OF ABBREVIATIONS

FPP	floating-point processor
LPHDR	low precision high dynamic range
PE	processing element
POSA	person of ordinary skill in the art
SIMD	single instruction multiple data

STATEMENT OF THE ISSUES

1. Whether substantial evidence supports the Board's finding that a POSA would not have been motivated to incorporate multiple Dockser FPPs as MacMillan's PEs where the costs associated with the complexity and size of Dockser FPPs would have defeated MacMillan's objective of avoiding increased costs associated with added complexity and size.

2. Whether substantial evidence supports the Board's finding that a POSA would not have been motivated to remove the traditional-precision capability of Dockser FPPs where Dockser and Tong teach the necessity of operating at traditional-precision with the ability to select a lower precision.

3. Whether, for certain claims, the Board's finding (that unpatentability has not been shown) should be affirmed on the alternate ground that the offered combination of prior art cannot meet the "exceeds" limitation.

INTRODUCTION

The Patent Trial and Appeal Board properly found that claims 3 and 23 of the U.S. Patent No. 10,416,961 (the "'961 Patent"), claims 3-8 of U.S. Patent No. 9,218,156 (the "'156 Patent"), and claims 3-20, 25, 34-61, 63, and 67-70 of U.S. Patent No. 8,407,273 (the "'273 Patent") (collectively, the "claims at issue") are not unpatentable as obvious. This Court should affirm that decision because substantial evidence supports the Board's factual findings that a POSA would not have been motivated to combine Google LLC's ("Google") references.

In its Brief, Google essentially argues that the Board ***could have*** found in Google's favor (in other words, Google argues that the record contains substantial evidence for its position). But that is not the test in this appeal. The test is whether substantial evidence supports the Board's factual findings—those findings lead inevitably to the Board's ultimate conclusion as to nonobviousness. A thorough review of the proceedings before the Board confirms that the Board's decision is well-supported.

Google filed three substantively similar petitions for *inter partes review* against the '961, '156, and '273 Patents. After discovery, briefing, and oral argument, the Board determined that Google had not shown that the claims at issue were obvious. Most of the claims at issue contain a limitation requiring that the number of low precision, high dynamic range (LPHDR) execution units exceeds (in

some cases, by a stated amount) the number of execution units adapted to perform traditional, high-precision multiplication in a particular device.

To try to meet this limitation, Google relied on two references—Dockser and MacMillan—in one ground, and the combination of Dockser, MacMillan, and Tong in a separate ground. Dockser discloses a processor, which it calls an FPP, that performs traditional-precision operations in default mode, but includes additional circuitry to allow for the selection of lower-precision operations as well. MacMillan is directed toward a low-cost supercomputing device that avoids complexity to minimize cost; MacMillan explains that its invention is based on using a multitude of simple processing elements, called PEs, that can be incorporated into one device.

In the IPRs, Google argued that the claims at issue are obvious because a POSA would be motivated to incorporate a large number of Dockser FPPs into a MacMillan device in place of MacMillan's PEs.

The Board determined that the claims at issue were not obvious. The Board made the factual findings that a POSA would not have been motivated to include Dockser FPPs into a MacMillan device because both Singular's and Google's experts agreed that a Dockser FPP is more complex than a conventional execution unit. Given MacMillan's express teachings that complexity and material costs should be avoided, a POSA would not have replaced MacMillan's low-cost PEs with Dockser FPPs.

Substantial evidence, including the disclosures of Dockser, MacMillan, and Tong, as well as the testimony of both experts, supports the Board's decision. This Court should, therefore, affirm.

STATEMENT OF THE CASE

I. THE CHALLENGED PATENTS

The patents at issue are directly related and share a common specification. Each claims priority, through parent and grandparent applications, to U.S. Provisional Patent Application No. 61/218,691, filed on Jun. 19, 2009.

A. The '961 Patent

The '961 Patent, entitled "Processing with Compact Arithmetic Processing Element," issued on September 17, 2019. Appx261.

The patent recognized that even though then-modern conventional microprocessors contained hundreds of millions of transistors, they could perform only a handful of operations per clock cycle. Appx274, 1:40-58. The patent explained that a large portion of this inefficiency comes from using transistor-intensive (*i.e.*, large) traditional-precision[1] execution units (also referred to as "processing elements" or "PEs") to perform arithmetic operations:

> As described above, today's CPU chips make inefficient use of their transistors . . . they deliver great precision, performing exact arithmetic . . . standardized arithmetic with 32 and 64 bit floating point numbers. Many applications need this kind of precision. As a result, conventional CPUs typically are designed to provide such precision, using on the order of a million transistors to implement the arithmetic operations.

[1] "Traditional precision" is defined herein as floating-point arithmetic with a word length of 32 or more bits, *see* Appx287 at 28:15-16.

Appx274-275, 2:62-3:10; Appx4963-4964, ¶ 39. *See also* Appx7886-7887, ¶ 39; Appx9664-9665, ¶ 39.

However, the patent described that such inefficient, traditional-precision execution units were not necessary for all applications:

> There are many economically important applications, however, which are not especially sensitive to precision and that would greatly benefit, in the form of application performance per transistor, from the ability to draw upon a far greater fraction of the computing power inherent in those million transistors. Current architectures for general purpose computing fail to deliver this power.

Appx275, 3:11-17; Appx4964, ¶ 40; Appx7887, ¶ 40; Appx9665, ¶ 40.

The '961 patent is thus directed away from prior art computers based on traditional-precision execution units that take up space and are wasteful of transistors. Appx4964-4965, ¶ 41; Appx7887-7888, ¶ 41; Appx9665-9666, ¶ 41.

As the patent further explains, "embodiments of the present invention may be implemented as any kind of machine which uses LPHDR arithmetic processing elements to provide computing using a small amount of resources (*e.g.*, transistors or volume) compared with traditional architectures." Appx277, 8:1-5; Appx4964-4965, ¶ 41; Appx7887-7888, ¶ 41; Appx9665-9666, ¶ 41. As a result, "[b]ecause LPHDR processing elements are relatively small, a single processor or other device may include a very large number of LPHDR processing elements, adapted to operate

in parallel with each other." Appx276, 6:47-52. The claims at issue in this appeal recite the presence of large numbers of LPHDR execution units within the claimed device.

By using a "very large number" of relatively small LPHDR execution units in parallel, computer systems can achieve significantly better performance than prior art systems for applications that are "not especially sensitive to precision." Appx275, 3:11-17; Appx4964, ¶ 40; Appx7887, ¶ 40; Appx9665, ¶ 40.

> PEs implemented according to certain embodiments of the present invention may be relatively small for PEs that can do arithmetic. This means that there are many PEs per unit of resource (e.g., transistor, area, volume), which in turn means that there is a large amount of arithmetic computational power per unit of resource. This enables larger problems to be solved with a given amount of resource than does traditional computer designs.

Appx285, 23:57-64.

In particular, the claimed systems "might perform tens of thousands of arithmetic operations per cycle, as opposed to hundreds in a conventional GPU or a handful in a conventional multicore CPU." *Id.*, 23:64-24:2; Appx4965, ¶ 42; Appx7888, ¶ 42; Appx9666, ¶ 42.

In addition, the '961 Patent also teaches computer systems in which the number of LPHDR execution units exceeds the number of traditional-precision execution units:

> For certain devices . . . according [to] the present invention, the number of LPHDR arithmetic elements in the device (e.g., computer or processor or other device) exceeds the number, possibly zero, of arithmetic elements in the device which are designed to perform high dynamic range arithmetic of traditional precision (that is, floating point arithmetic with a word length of 32 or more bits).

Appx287, 28:9-16; Appx4966, ¶ 43; Appx7889, ¶ 43; Appx9667, ¶ 43.

Such heterogenous architectures allow applications to make use of both traditional-precision and reduced-precision operations while retaining massive scale and parallelism in the form of a large number of LPHDR execution units (*see* for example Appx282, 18:57-64).

The increased level of compute parallelism and scale in such computer systems is necessarily achieved at the cost of precision—a significant percentage of the high dynamic range floating-point operations performed by the LPHDR execution units of the claimed device produce outputs that differ significantly from the result of an exact mathematical calculation of those same operations. By sacrificing precision for increased parallelism/scale, the patent enables new computer architectures that feature significant performance gains per unit of resource over prior art computer architectures. Appx4966, ¶ 44; Appx7889, ¶ 44; Appx9667, ¶ 44.

> To conclude, modern digital computing systems provide high precision arithmetic, but that precision is costly. A modern

> double precision floating point multiplier may require on the order of a million transistors, even though only a handful of transistors is required to perform a low precision multiplication. Despite the common belief among those having ordinary skill in the art that modern applications require high precision processing, in fact a variety of useful algorithms function adequately at much lower precision. As a result, such algorithms may be performed by processors or other devices implemented according to embodiments of the present invention, which come closer to achieving the goal of using a few transistors to multiply… thus enabling massively parallel arithmetic computation to be performed with relatively small amounts of physical resources (such as a single silicon chip).

Appx277, 7:8-24.

B. The ’156 Patent

The ’156 Patent (Appx291) shares a specification with the ’961 Patent.[2] and claims different variations of the invention.

C. The ’273 Patent

The ’273 Patent (Appx320) also shares a specification with the ’961 Patent and claims different variations of the invention.

[2] Citations in this Brief are to the ’961 Patent, as the ’156 and ’273 Patents have an identical disclosure.

II. OVERVIEW OF THE CITED REFERENCES

A. Dockser

U.S. Pat. Publ. 2007/0203967 ("Dockser") discloses the use of a traditional-precision floating point processor (FPP) that can selectably reduce precision in order to reduce its power draw. *See, e.g.*, Appx1472, ¶ [0026]; Appx4967, ¶ 45; Appx7890, ¶ 45; Appx9668, ¶ 45.

Dockser's teaching of an execution unit that is always capable of operating in traditional-precision mode (its default), but that could occasionally operate in a reduced-precision mode, was motivated by the prevalent view in the prior art that while low-precision operation might be acceptable for "certain applications," for a general-purpose processor traditional-precision capability is "needed." Appx1470; Appx1471, ¶¶ [0003, 0018], claims 1, 8, 15, 20; Appx4967, ¶ 46; Appx7890, ¶ 46; Appx9668, ¶ 46.

The components of Dockser's FPP are shown in Figure 1:

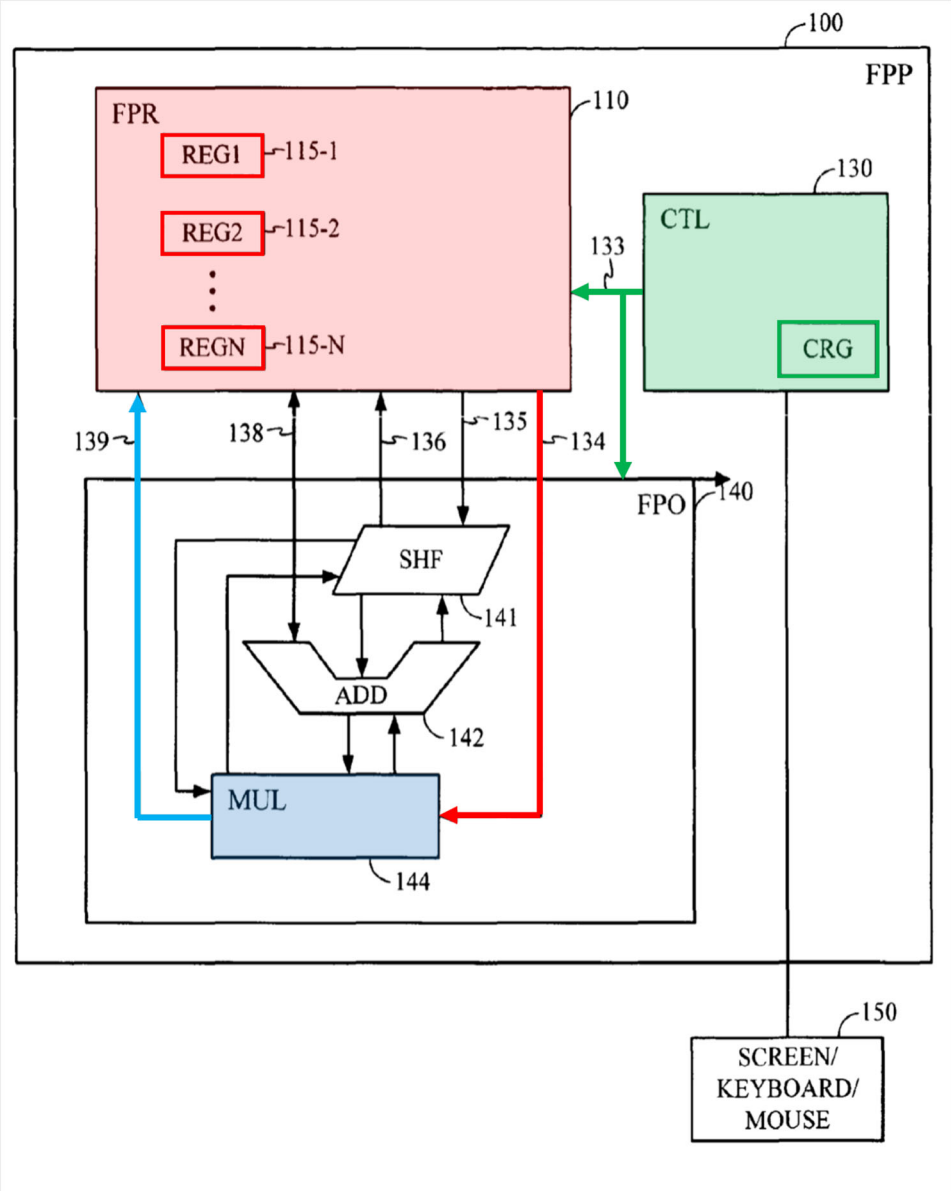

Dockser, Fig. 1 (Annotated)

Dockser's FPP has a register file (FPR 110, shaded in **red**) comprising registers (red boxes) that hold IEEE 32-bit traditional-precision values. It also includes a controller (CTL 130, shaded in **green**) with a control register (**green** box) that can store "subprecision select bits" corresponding to the desired level of precision. Appx1466, Fig. 1; *see also* Appx1471, ¶¶ [0017-18]; Appx4967-4969,

¶ 47; Appx7890-7892, ¶ 47; Appx9668-9670, ¶ 47. The FPP performs arithmetic operations (*e.g.*, addition and multiplication) at traditional-precision or at the desired sub-precision using values stored in the register file as operands. *See* Appx1471-1472, ¶ [0019]; Appx4967-4969, ¶ 47; Appx7890-7892, ¶ 47; Appx9668-9670, ¶ 47.

If no sub-precision is selected, Dockser's FPP performs arithmetic at "maximum precision" (*i.e.*, 32-bit IEEE traditional precision). *See* Appx1474-1475, claims 1, 8, 15, 20; *see also* Appx1471, ¶ [0018] ("the floating-point controller 130 *may* be used to select the subprecision of the floating-point operations")[3]; Appx4969, ¶ 48; Appx7892, ¶ 48; Appx9670, ¶ 48. In reduced-precision modes, Dockser's FPP reduces the mantissa size by powering-down some of the cells in the register file that store the mantissa, and also powering-down parts of the arithmetic logic circuits (contained, *e.g.*, within the multiplier, shaded in **blue** above) that are not needed for the selected sub-precision. Appx1472, ¶¶ [0026-27]; Appx4969, ¶ 48; Appx7892, ¶ 48; Appx9670, ¶ 48.

Unlike the '961 Patent, Dockser does not teach a system that includes multiple execution units operating in parallel. Appx4750, 71:13-21; Appx4969, ¶ 49; Appx7892, ¶ 49; Appx9670, ¶ 49. Rather, Dockser discloses a single floating-point execution unit in isolation. *Id.* The absence in Dockser of any parallel computing teaching is not surprising given its selectable low-precision unit is designed to

[3] All emphasis in this Brief is added unless otherwise indicated.

conserve power in "battery operated devices where power comes at a premium, such as wireless telephones, personal digital assistants (PDA), laptops, game consoles, pagers, and cameras." Appx1470, ¶ [0003]; Appx4969, ¶ 49; Appx7892, ¶ 49; Appx9670, ¶ 49.

Because Dockser's FPP is capable of traditional precision, it would have at least as many transistors and take up at least as much space as a conventional traditional-precision execution unit, even when operating in a reduced-precision mode. Appx4750, 71:6-12; Appx4743, 43:21-45:9; Appx4970, ¶¶ 50-51; Appx7893, ¶¶ 50-51; Appx9671, ¶¶ 50-51. Because Dockser's FPP uses a traditional-precision execution unit and includes additional control circuits for selecting reduced-precision modes, it is likely even larger than a conventional traditional-precision execution unit, making it unsuitable for scaling the compute power of a conventional computer. *Id.*

Dockser thus does not mention increasing compute scale and parallelism and does not teach a computer system comprised of a larger (let alone much larger) number of low-precision execution units than traditional-precision execution units.

B. Tong

"Reducing Power by Optimizing the Necessary Precision/Range of Floating-Point Arithmetic" by Jonathan Ying Fai Tong, *et al.,* Appx1476, teaches that using lower-precision arithmetic can reduce power consumption. *See* Appx1483. Tong

however, like Dockser, is also based on the then commonly held belief that computer systems *must* support a traditional-precision mode because some applications can only be performed using traditional-precision arithmetic: "[e]ven though we may be able to assume that most of our operands can be computed successfully in limited precision, it appears inevitable that some fraction of our operands will require full IEEE-standard precision." *Id.*; Appx4970, ¶ 52; Appx7893, ¶ 52; Appx9671, ¶ 52.

Importantly, like Dockser, Tong is focused on reducing power consumption and does not teach parallel processing computer systems that include large numbers of execution units that operate simultaneously, let alone systems having much larger numbers of low-precision units than traditional-precision execution units. Appx4971, ¶ 53; Appx7894, ¶ 53; Appx9672, ¶ 53.

Each low-precision unit disclosed by Tong (floating-point unit, or FP unit) is either as large as a traditional-precision execution unit (because it includes complex circuitry to combine the results of multiple sequential operations into one traditional-precision result) or paired with a traditional-precision execution unit on a 1:1 basis (Appx1485). Appx4971, ¶ 53; Appx7894, ¶ 53; Appx9672, ¶ 53. For example, Tong teaches "simply including both full and reduced precision FP units and using appropriate sleep-mode circuit techniques to shut down the unused unit." Appx1485; Appx4971, ¶ 54; Appx7894, ¶ 54; Appx9672, ¶ 54. In this approach, each low-precision unit is paired with a traditional-precision execution unit in a 1:1

ratio, with the resulting pair requiring more physical space than a traditional-precision execution unit alone. *Id.* Indeed, this approach is presented as an option in situations where silicon real-estate is not at a premium. *See* Appx1485 ("Given the decreasing cost of silicon area . . . ."); Appx4971, ¶ 54; Appx7894, ¶ 54; Appx9672, ¶ 54.

As with Dockser, Tong is focused on power-savings, does not even mention increasing compute scale and parallelism, and does not teach a computer comprised of a larger (let alone much larger) number of low-precision units than traditional-precision execution units. Appx4972, ¶ 56; Appx7895, ¶ 56; Appx9673, ¶ 56.

C. MacMillan

U.S. Patent No. 5,689,677 ("MacMillan"), Appx1490, is directed to a computer system that includes a host processor and "a plurality of processing elements." Appx1502, 12:39. MacMillan does not describe the capabilities of each "processing element" (PE) in detail, noting only briefly that "[i]nteger and floating point accelerators could be included in each PE." *Id.*, 12:55-56; Appx4972-4973, ¶ 57; Appx7895-7896, ¶ 57; Appx9673-9674, ¶ 57.

Indeed, MacMillan is silent as to arithmetic precision. Appx4973, ¶ 58; Appx7896, ¶ 58; Appx9674, ¶ 58. It simply teaches that each PE can "perform atomic operations on data values up to *32 bits wide.*" Appx1500, 7:8-9; Appx4973, ¶ 58; Appx7896, ¶ 58; Appx9674, ¶ 58. Rather, MacMillan is directed towards

providing supercomputer capabilities, without incurring substantial cost or unnecessary complexity: "[i]f supercomputing performance could be achieved in low cost computers, such as personal computers, this could dramatically expand the market for personal computers." Appx1497, 1:21-24. MacMillan explains the use of conventional SIMD systems results in higher costs. *Id.*, 2:57–61 ("hav[ing] a large register set on each PE . . . increases die area per PE, resulting in higher costs per PE"); *id.*, 2:62–65 ("Adding pins can reduce the PE-to-memory bottleneck, but leads to increased packaging costs. It may also require tighter geometries or increased numbers of layers on printed circuits boards, further increasing costs."); Appx1498, 3:2–3 ("Adding output buffers to drive increased pin counts also increases power dissipation and hence power supply capacity and cost"); *id.*, 4:39–40 ("The PE array[] disk may be limited in size due to technology or cost factors.")

MacMillan therefore instructs that "[t]o meet the cost objectives, the SIMD capabilities should not add significant complexity to the architecture of a computer system for personal use." Appx1499, 5:42-44; *see also* 5:58-59. To achieve this goal, MacMillan explains that its invention should use conventional components, and that PEs should not include additional registers. Appx1497, 2:56-61 (explaining drawbacks of a PE with "a large register set"); Appx1499, 6:24-26, 34-36 (requiring use of conventional components). MacMillan's disclosed architecture is designed to overcome the limitations of more complex SIMD systems, to allow for high

performance at low cost. *Id.*, 6:41–45; Appx1504, 16:14–27; *see also* Appx1499, 5:38-61, 6:1-10.

Unlike Dockser and Tong, MacMillan is focused on the provision of low cost supercomputing using simplified components, and is not specifically tailored for the low requirements in power portable or mobile devices. *See* Appx1498, 3:4-6; Appx4973, ¶ 58; Appx7896, ¶ 58; Appx9674, ¶ 58.

III. THE IPR PROCEEDINGS

Google filed three separate IPR Petitions – IPR2021-00155, IPR2021-00165, and IPR2021-00179 -- challenging claims 1-5, 10, 13-14, 21, and 23-25 of the '961 Patent; claims 1-8, 16, and 33 of the '156 Patent; and claims 1-26, 28, 32-61, 63, and 67-70 of the '273 Patent, respectively. Appx652; Appx6359; Appx8334. The Board instituted trial in each proceeding on May 14, 2021. Appx3972; Appx7661; Appx9442.

The claims could be divided into two groups, with one group directed to devices with "at least one LPHDR execution unit," and the other group directed to devices with multiple LPHDR execution units. The Board found Google had proved the "at least one LPHDR execution unit" group to have been obvious but had failed to sustain its burden as to the other "multiple LPHDR execution units" group. Singular appealed the decisions as to the "at least one LPHDR execution unit" group

but has since withdrawn those appeals. Thus, the remaining appeals concern only the claims directed to devices with multiple LPHDR execution units.

The relevant grounds to these appeals are Grounds 3 and 4 of each Petition. In Ground 3, Google argued that a POSA would replace MacMillan's PEs with Dockser FPPs to realize some unspecified decrease in power consumption. In Ground 4, Google argued that either: i) a POSA would combine Dockser with MacMillan and use Tong to set precision levels that meet the claim limitations as to precision; or ii) a POSA would physically remove Dockser's traditional-precision capability to make the Dockser FPP no longer "adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide."

The Board issued Final Written Decisions in each case. Appx1; Appx89; Appx174. The decisions are similar in relevant part. In sum, the Board found obvious the claims reciting "at-least-one" LPHDR execution unit but determined that claims reciting multiple LPHDR execution units were not unpatentable as obvious. In particular, the Board found that a POSA would not have been motivated to combine Dockser and MacMillan and, further, that a POSA would not have been motivated to combine Dockser, MacMillan, and Tong.

A. The Board's Determination that a POSA Would Not Combine Dockser and MacMillan

The Board determined that a POSA would not combine Dockser with MacMillan. Appx72; Appx160; Appx245. As described above, Dockser requires

additional circuitry and complexity to allow selection of various arithmetic precisions. Appx72. MacMillan, in contrast, is "heavily focused on avoiding increased costs associated with added complexity." *Id.* The Board noted that MacMillan described "numerous limitations of known SIMD designs at the time, many of which related to increased costs from added complexity." *Id*. Indeed, MacMillan discloses an architecture "with multiple, less costly PEs" that provides a low-cost solution. Appx73; Appx161; Appx246.

With those key principles of both references in mind, the Board determined that a POSA would not be motivated to replace MacMillan's PEs with Dockser's FPPs. Appx72-77; Appx156-161; Appx242-247. The testimony of both experts established that a Dockser FPP required more components, took up more space, and was thus more complex and costly than a MacMillan PE . *Id.* Given MacMillan's emphasis on avoiding increased costs associated with added complexity and size, a POSA would not use a Dockser FPP in MacMillan. *Id.*

Accordingly, the Board found that claim 3 of the '961 Patent, claims 3-8 of the '156 Patent, and claims 3-20, 25, 36-61, and 63 of the '273 Patent were not shown to be unpatentable as obvious under Ground 3.

B. The Board's Determination that a POSA Would Not Combine Dockser, Tong, and MacMillan

As to Ground 4, the Board found that a POSA would not be motivated to combine Dockser and MacMillan for the same reasons expressed regarding Ground

3. Appx78; Appx162; Appx249. With respect to Google's proposal that a POSA would physically remove Dockser's traditional-precision capability, the Board found that both Dockser and Tong teach away from such a modification. Appx80-83; Appx164-167; Appx251-254. A POSA would not be motivated to modify Dockser or Tong against those references' express statements. A POSA also would not be motivated to make such modifications against MacMillan's own express statements. Even a Dockser FPP, if modified as suggested by Google to no longer have traditional-precision capability, would retain some precision selection ability (*see, e.g.*, Appx5137-5138; Appx8059-8060; Appx9838-9839) which means it would thus introduce additional complexity against MacMillan's express teachings. Appx82; Appx166-167; Appx254.

Accordingly, the Board found that claims 3 and 23 of the '961 Patent, claims 3-8 of the '156 Patent, and claims 3-20, 25, 34-61, 63, and 67-70 of the '273 Patent were not shown to be unpatentable as obvious under Ground 4.

IV. PROCEEDINGS IN THIS COURT

On June 2, 2022, Google filed its Notices of Appeal in the PTAB. Appx6247; Appx8229; Appx10019. On July 12, 2022, Singular filed its Notices of Appeal. Appx6351; Appx8326; Appx10118 as to the decisions invalidating the "at-least-one" LPHDR execution unit claims. Appx6351; Appx8326; Appx10118. On December 16, 2022, Singular filed an unopposed motion to dismiss its cross-appeals.

Accordingly, the remaining issues in this case are whether the PTAB correctly determined that the multiple LPHDR execution unit claims -- claims 3 and 23 of the '961 Patent, claims 3-8 of the '156 Patent, and claims 3-20, 25, 34-61, 63, and 67-70 of the '273 Patent (*i.e.*, the claims at issue) -- are not unpatentable as obvious.

SUMMARY OF THE ARGUMENT

The Court should affirm the Board's decision. Substantial evidence supports the Board's factual findings that a POSA would not have been motivated to combine Dockser and MacMillan. During the IPR, Singular presented testimony from both its own and Google's experts that a Dockser FPP—which includes additional circuitry to select precision—requires more circuitry, and thus more complexity, space, and cost, than a conventional floating-point unit. As Singular's expert explained, and as Google did not challenge, MacMillan teaches away from introducing additional complexity or cost in its PEs. Given the Board's factual findings are supported by substantial evidence, the Court should affirm the Board's ultimate determination of non-obviousness.

Alternatively, the Court should affirm as to most of the claims at issue because Google could not show that its proposed combination disclosed or rendered obvious the claim limitations requiring that the number of LPHDR execution units exceeds the number of execution units "adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide." Dockser's FPP always includes traditional-precision capability. This means that even assuming a motivation to combine had existed, each substitution of a Dockser FPP for a MacMillan PE would add a unit that is both an LPHDR execution unit and an execution unit "adapted to execute at least the operation of multiplication on floating

point numbers that are at least 32 bits wide." Accordingly, a Dockser-MacMillan or Dockser-MacMillan-Tong combination cannot meet this "exceeds" claim limitation.

STANDARD OF REVIEW

This Court reviews the Board's legal determinations *de novo* but reviews the Board's factual findings underlying those determinations for substantial evidence. *Mylan Pharms. Inc. v. Merck Sharp & Dohme Corp.*, 50 F.4th 147, 152 (Fed. Cir. 2022) (citations omitted). Obviousness is a mixed question of law and fact, and the Court reviews the Board's ultimate obviousness determination *de novo* and underlying fact-findings for substantial evidence. *Id.*

Substantial evidence means "such relevant evidence as a reasonable mind might accept as adequate to support a conclusion." *Consol. Edison Co. v. NLRB*, 305 U.S. 197, 229 (1938); *In re Gartside*, 203 F.3d 1305, 1312 (Fed. Cir. 2000).

The presence or absence of a motivation to combine references in an obviousness determination is a pure question of fact. *Id.* at 1316; *In re Dembiczak*, 175 F.3d 994, 998 (Fed. Cir. 1999).

The Board's findings regarding the scope and content of the prior art are reviewed for substantial evidence. *Polaris Innovations, Ltd. v. Brent*, 48 F.4th 1365, 1380 (Fed. Cir. 2022); *Amazon.com, Inc. v. Barnesandnoble.com, Inc.*, 239 F.3d 1343, 1358 (Fed. Cir. 2001) ("[W]hat a reference teaches is a question of fact.").

In an IPR, the burden of persuasion is on the petitioner to prove unpatentability by a preponderance of the evidence and that burden never shifts to the patentee. *Google LLC v. IPA Tech. Inc.*, 34 F.4th 1081, 1085 (Fed. Cir. 2022).

ARGUMENT

V. THE BOARD'S FINDING THAT A POSA WOULD NOT HAVE BEEN MOTIVATED TO COMBINE DOCKSER WITH MACMILLAN BECAUSE IT WOULD HAVE DESTROYED MACMILLAN'S STATED OBJECTIVE OF AVOIDING INCREASED COST WAS SUPPORTED BY SUBSTANTIAL EVIDENCE

The Court should affirm with respect to Ground 3 of the Petition because substantial evidence supported the Board's finding that a POSA would not have been motivated to combine Dockser and MacMillan. *Outdry Techs. Corp. v. Geox S.p.A.*, 859 F.3d 1364, 1368 (Fed. Cir. 2017) ("The Board's motivation to combine finding is reviewed for substantial evidence"); *Redline Detection, LLC v. Star Envirotech, Inc.*, 811 F.3d 435, 449 (Fed. Cir. 2015) ("A finding is supported by substantial evidence if a reasonable mind might accept the evidence to support the finding") (quoting *K/S Himpp v. Hear-Wear Techs., LLC*, 751 F.3d 1362, 1364 (Fed. Cir. 2014)).

The Board considered the disclosures of Dockser and MacMillan, as well as the parties' submitted evidence, including the testimony of Mr. Goodin and Dr. Khatri. Appx72; Appx156-157; Appx242-243. Based on its evaluation of that evidence, the Board determined that Google had not shown a motivation to use Dockser's FPP in place of the PEs within MacMillan.

A. Dockser and MacMillan Were Directed to Different Problems

In its Petition, Google had argued that a POSA would have been motivated to combine Dockser and MacMillan "to increase performance speed as MacMillan teaches while lowering power consumption as Dockser teaches." Appx715; Appx6419-6420; Appx8392-8393. Google's proposed motivation ignored the fact that the disclosure of MacMillan focused on providing supercomputer capabilities *while* avoiding complexity that would increase costs. As the Board noted, MacMillan criticizes existing designs due to their increased costs, such as from including additional registers or other complexity. *Id.*

The inventive concept of MacMillan is to "allow supercomputing performance to be provided in a low cost computer system for personal use, dramatically expanding the potential market for systems with supercomputer performance. To meet the cost objectives, the SIMD capabilities should not add significant complexity to the architecture of a computer system for personal use. The present invention addresses the above needs." Appx1499, 5:39-45.

MacMillan's invention addresses the needs of supercomputer performance at reduced cost by providing an architecture that enables less complex and smaller PEs to work together. MacMillan's PEs should be simple and small. For example, MacMillan states that "the SIMD capabilities should not add significant complexity to the architecture of a computer system" and that "hav[ing] a large register set on

each PE . . . increases die area per PE, resulting in higher costs per PE." *Id.*, 5:42-44; Appx1497, 2:57-61.

In contrast to MacMillan's directive to reduce complexity, Dockser's FPP is more complex and larger than a traditional-precision execution unit because of its need for additional control circuitry to implement selectable sub-precision modes. In the IPR, both experts agreed that a Dockser FPP was more complex and larger than a traditional execution unit. Appx4992-4993, ¶ 103; Appx7915-7916, ¶ 103; Appx9693-9694, ¶ 103; Appx4750, 70:6-12 (explaining that Dockser does not reduce space compared to conventional units), Appx4743, 43:21-45:9 (additional circuitry required to reduce precision). As the Board concluded from the evidence, inserting Dockser's complex execution units (the variable precision FPPs) into MacMillan's system would conflict with MacMillan's admonition against increasing complexity in order to minimize costs. Appx4993, ¶ 104; Appx7916, ¶ 104; Appx9694, ¶ 104. A POSA would thus not have been motivated to make a combination that would have defeated MacMillan's objective. *Trivascular, Inc. v. Samuels*, 812 F.3d 1056, 1068 (Fed. Cir. 2016); *Chemours Co. FC, LLC v. Daikin Indus., Ltd.*, 4 F.4th 1370, 1376-77 (Fed. Cir. 2021).

B. The Board's Finding Was Based on Substantial Evidence as to These Differences.

The Board examined Dockser and MacMillan's disclosures and the testimony of Dr. Khatri and Mr. Goodin. Appx72-77; Appx160-165; Appx245-250. The

Board noted MacMillan's teachings that the PEs should be low cost and have low complexity, allowing for many low-cost processors in a cost-effective device. Appx76; Appx164; Appx249. Moreover, the Board considered Mr. Goodin's testimony admitting that a Dockser FPP is of greater complexity than a standard processor. *Id.* Finally, the Board considered Dr. Khatri's uncontroverted testimony on this issue. *Id.*

The Board's finding that a POSA would not have been motivated to combine the teachings of these references was explicitly based on the evidence of the differences between them. A reasonable mind would accept the evidence supporting the Board's finding. Accordingly, the Court should affirm.

C. The Board Properly Considered that Combining Dockser with MacMillan Would Increase Cost Because of Added Complexity

Google primarily argues that the Board erred in considering the increased complexity and cost that would result from inserting Dockser's FPP into MacMillan's system in its motivation to combine analysis. Blue Br. at 36-42. Google misinterprets both the law and the Board's decision.

Singular argued that the increased complexity and cost of inserting many Dockser FPPs into a MacMillan system would discourage a POSA from combining the two references. Google did not present any contrary argument to the Board on this point during the IPR. Indeed, the entirety of Google's discussion before the

Board about the effect of increased complexity and cost on this proposed combination was a single sentence in its Reply. Appx5051-5052; Appx7973-7974; Appx9752-9753 (arguing that increased costs "would be legally irrelevant because whether a combination would not be made by businessmen for economic reasons is not probative of nonobviousness" (internal quotations and citations omitted)).

In any event, Google's cases. *In re Farrenkopf*, 713 F.2d 714 (Fed. Cir. 1983), and the other cited cases considered whether a combination's additional economic cost standing alone bars a finding of nonobviousness where a POSA would otherwise have had sufficient reason to build the more expensive combination. Here, by comparison, MacMillan bases its inventive purpose on minimizing costs, and Dockser has a design that would negate that purpose. The inventive concept in MacMillan was a low-complexity, smaller-computer system that adds parallel SIMD processing "inexpensively to existing system architectures." Appx1499, 6:2-3. Introducing additional complexity, size, and therefore cost, in the form of many Dockser FPPs, would defeat that purpose. The Board's conclusion thus was based on the incompatibility of Docker's more complex FPPs with MacMillan's objective of reducing cost by reducing complexity. Google does not challenge that MacMillan teaches providing low-cost supercomputer capabilities using simple PEs. Since Dockser's FPP is indisputably more complex and thus increases cost relative to a MacMillan PE, the Board properly found that a POSA would not have been

motivated to combine them. Google cannot shoehorn this case, in which the Board's motivation to combine analysis turns on the incompatibility of the design of one reference with the stated objective of another reference, into its cited line of cases where the motivation to combine analysis merely turns on economic cost of combining otherwise compatible references.

Farrenkopf held that a lone statement in a reference that "the use of inhibitors can be costly" was not sufficient to defeat a motivation to combine when the secondary references taught other advantages: "That a given combination would not be made by businessmen for economic reasons does not mean that persons skilled in the art would not make the combination because of some technological incompatibility." *Farrenkopf*, 713 F.2d at 718. In the instant case, by contrast, the reason the Board did not find a motivation to combine was because of such a "technological incompatibility."

Indeed, a "technological incompatibility" would have been particularly relevant to a POSA here because MacMillan is directed to *reducing* the complexity of existing supercomputing systems in order to reduce costs, while incorporating Dockser's FPPs would increase complexity and therefore costs. Here, reducing complexity and cost is not merely a benefit important only to "businessmen for economic reasons," as Google argues. To the contrary, reducing cost and

complexity was the expressly stated technological objective of the MacMillan system.

Similarly, *Orthopedic Equip. Co., Inc. v. United States*, 702 F.2d 1005, 1012-13 (Fed. Cir. 1983), held that motivation to combine would not be defeated just because the resulting device was not attractive to business. Here, MacMillan expressed a stated objective—lower cost through lower complexity—that according to Google's own expert was flouted by Dockser's increased cost and complexity.

Further, *In re Nilssen*, 837 F.2d 1098 (Fed. Cir. 1987), lacked the salient fact that one reference in the proposed combination flouted a stated objective of another reference in the combination. *Friskit, Inc. v. Real Networks*, Inc., 306 F. App'x 610, 617-18 (Fed. Cir. 2009), is even less relevant. There, the Court held only that long-felt need could not be shown by purely economic evidence. *In re Magna Elecs., Inc.*, 611 F. App'x 969, 973 (Fed. Cir. 2015), came to a similar conclusion as to skepticism of experts. Again, none of these cases featured an express teaching in one reference that was flouted by the combination's other reference.

Finally, Google cites *Grit Energy Sols., LLC v. Oren Techs., LLC*, 957 F.3d 1309, 1323 (Fed. Cir. 2020), where the petitioner had argued to the Board a POSA would make a given combination because that combination "would reduce the cost of the system." The Board, however, determined "that the proposed [combination] would have increased the cost of the system," and the court then observed that

determination "might cut against the force of [petitioner's] argument." The court obviously thought that costs were relevant to a motivation-to-combine analysis because it then remanded for a more detailed explanation as to how the combination "would result in a more expensive system." *Id.*, 1324. Because the case suggests increased costs *can* be a factor to consider in the motivation to combine analysis, *Grit Energy Sols.* does not support Google's argument.

D. Substantial Evidence Supports the Board's Finding that Adding Dockser to MacMillan Would Increase Macmillan's Complexity, Size, and Therefore Cost

In the IPR, as discussed above, Singular provided substantial evidence that using Dockser's FPP in MacMillan would increase complexity, size, and therefore cost. The Board cited the disclosures of Dockser and MacMillan, expert testimony from Dr. Khatri, and cross-examination testimony from Mr. Goodin. Appx72; Appx156; Appx242. Google now argues this evidence is not sufficient to establish that adding precision control circuitry increases Dockser's size, complexity, and cost. Blue Br., pp. 43-48.

However, Google had the burden of persuasion to prove unpatentability by a preponderance of the evidence, including establishing that no such increase would have occurred. *IPA Techs.*, 34 F. 4th at 1085. Google did not meet that burden. In any event, no analysis over and above the plainly stated teachings of the combination's references is required. Google cites no law or case holding otherwise.

Even if such an analysis were required, Google itself provides it in tacitly admitting that Dockser's FPP is larger, more complex, and costlier than a traditional-precision execution unit. Blue Br. at 44. There is no dispute that Dockser shows additional circuitry used solely for its precision selection mechanism.

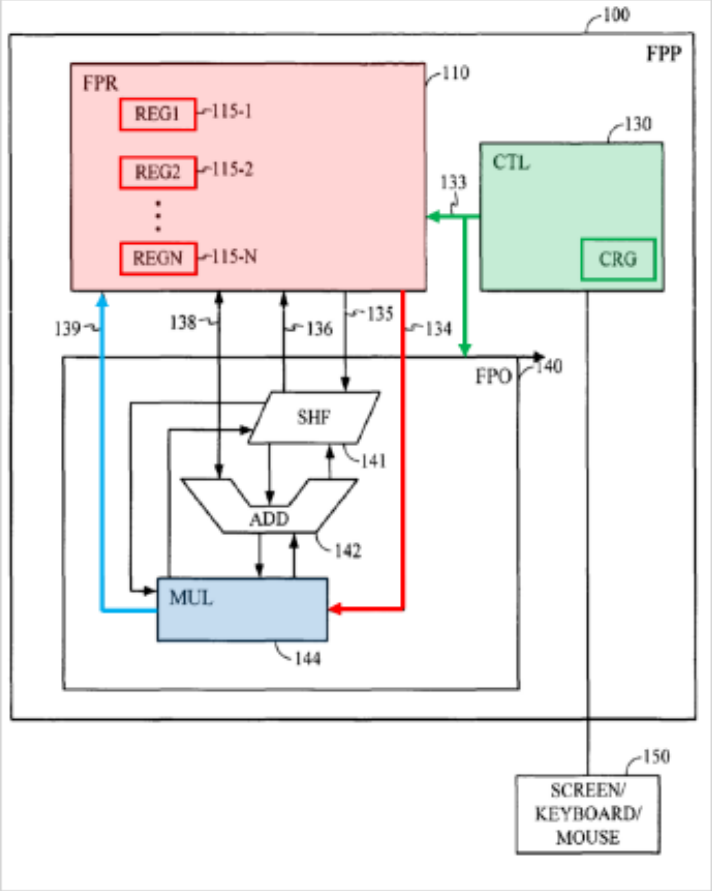

Dockser, Fig. 1 (Annotated)

Dockser Fig. 1 shows, at the least, the green shaded area above, which is used solely for its precision selection. Appx1466; *see also* Appx1471, ¶¶ [0017-18]; Appx4967-4969, ¶ 47; Appx7890-7892, ¶ 47; Appx9668-9670, ¶ 47. Given that the remainder of Dockser is implemented using conventional components, the additional precision selection logic—additional circuitry—would increase size,

complexity, and costs, as compared to the standard precision execution unit ("PE") of Macmillan. Appx74-75; Appx158-159; Appx244-245; Appx1471, ¶¶ [0017-18]; Appx4750, 71:6-12; Appx4743, 43:21-45:9; Appx4970, ¶¶ 50-51; Appx7893, ¶¶ 50-51; Appx9671, ¶¶ 50-51.

Similarly, Dr. Khatri's testimony supports the Board's finding. Dr. Khatri testified that the inclusion of Dockser's additional circuitry would increase size and complexity leading to increased costs. Appx4970, ¶¶ 50-51; Appx7893, ¶¶ 50-51; Appx9671, ¶¶ 50-51. Google faults Dr. Khatri's testimony for not including some unspecified "particular" analysis showing that additional circuitry increases size and cost. Blue Br. at 45. MacMillan itself, however, concludes as much: "[S]ome SIMD designs have a large register set on each PE, allowing temporary results to remain local to the PE, thereby reducing the number of loads and stores needed. However, this increases die area per PE, resulting in higher costs per PE." Appx1497, 2:57-61. *See also id.*, 2:63-65 (noting that increased number of layers increases costs); *id.* 2:67-3:2 (increasing die area "rapidly increase[s] costs"). Dr. Khatri's statement that circuitry costs money and takes up space, is self-evident to a POSA, and is supported by the MacMillan quotes above.

Finally, Google's own expert testimony supported the Board's finding that adding Dockser to MacMillan would increase MacMillan's complexity and therefore cost. The Board reasonably relied on that testimony. Appx71-72;

Appx156; Appx242. In this appeal, Google attempts to recast its expert's testimony. Blue Br. at 45-46. But Mr. Goodin testified unequivocally that a Dockser FPP would require circuitry to perform its precision selection. Appx4743, 43:21-45:9. Moreover, Mr. Goodin agreed that Dockser did not disclose any embodiment that would take up a smaller area than a standard 32-bit processor. Appx4750, 71:6-12.

Substantial evidence thus supports the Board's finding that Dockser introduces additional size, complexity, and therefore cost. *Provisur Techs., Inc. v. Weber, Inc.*, No. 21-1851, 2022 WL 17688071, at *5 (Fed. Dec. 15, 2022) (nonprecedential) ("The Board recognized that because of this increased complexity, a person of ordinary skill in the art would have lacked motivation to combine the references and modify them to obtain the reverse fill configuration. This testimony was substantial evidence for the Board's conclusion.").

Google raises new arguments not made during the IPR, unsupported by any evidence in the record. For example, Google now argues that MacMillan is only concerned with avoiding ***numerically specified*** amounts of complexity or increased size, as opposed to complexity and size generally. Blue Br. at 44, 47. Specifically, Google now argues that "MacMillan was concerned with avoiding the $100,000 price tag of prior-art supercomputers." *Id.* at 47. But Google did not present this argument to the Board. The sole argument Google presented regarding costs was a

lack of the unnecessary "specific cost analysis," discussed above. Appx5051-5052; Appx7973-7974; Appx9752-9753.

Similarly, Google has also waived its heretofore unpresented argument that a Dockser/MacMillan combination would result in a net cost *savings*. Blue Br. at 46. Google did not raise this argument before the Board, has not presented any evidence in its support, and thus has waived it before this Court.

E. The Board Did Not Commit a Legal Error by Finding MacMillan Teaches Away from Including Dockser

Google alternatively argues that the Board committed a legal error by failing to give the supposed benefits of Dockser and MacMillan "any weight in its analysis." Blue Br. at 48-51. But Google applies the wrong law to this question. The Board explicitly found that MacMillan teaches away from including a Dockser-type PE that would increase size, complexity, and costs. Appx76; Appx160; Appx246. "Under this court's precedent, a prior art reference is said to teach away from the claimed invention if a skilled artisan 'upon reading the reference, would be discouraged from following the path set out in the reference.'" *AstraZeneca AB v. Mylan Pharms. Inc.*, 19 F.4th 1325, 1337 (Fed. Cir. 2021). As the Board explained, "MacMillan's statements that its approach is designed to avoid increased costs from added complexity and size (and thereby allow for many low-cost processors in the device) would have dissuaded a person of ordinary skill in the art from" combining Dockser's and MacMillan's approaches. Appx76; Appx160; Appx246. No law

requires a comparison of the benefits of a combination when there is a teaching away; indeed, such an analysis (which would essentially compare the benefits of the claimed invention to the prior art) would inevitably be poisoned by hindsight. *See Chemours*, 4 F.4th at 1376 . In other words, Google would have this Court infer that a POSA would be motivated to ignore an express teaching away in a reference, solely because the resulting combination had the benefits of the invention asserted to be obvious. No invention could survive such a challenge and this Court has never applied such a test.

Moreover, this Court has repeatedly explained that there cannot be a motivation to combine references in a way that defeats their purpose. For example, in *Chemours*, this Court explained that there was insufficient motivation to "increase [a prior art reference's] melt flow rate to the claimed range," in part because the reference "includes numerous examples of processing techniques that are typically used to increase melt flow rate, which [the reference] cautions should *not* be used due to the risk of obtaining a broader molecular weight distribution." *Id.* at 1376-77. Similarly, in *TriVascular*, 812 F.3d at 1067–69, the Court found insufficient motivation to combine references "since [the petitioner's] proposed substitution would destroy the basic objective of the" prior art reference (*i.e.*, much like MacMillan's basic objective of reducing cost by reducing complexity and size). *See In re Gordon*, 733 F.2d 900, 902 (Fed. Cir. 1984) (finding no reason to modify a

prior art device where the modification would render the device "inoperable for its intended purpose"); *see also Henny Penny Corp. v. Frymaster LLC*, 938 F.3d 1324, 1332 (Fed. Cir. 2019) (affirming the Board's finding of no motivation to combine where the Board "credited [the patent owner's] expert's testimony that following [one reference's] method of diverting and cooling the oil in [another reference's] system would introduce 'additional plumbing and complexity'").

Google's cases on this issue do not support its position. For example, in *Intel Corp. v. Qualcomm Inc.*, 21 F.4th 784 (Fed. Cir. 2021), the issue was whether there was a motivation to take a switch discussed in Burgener and combine it by incorporation into the Der reference's system (which Burgener references as "prior art"). *Id.* at 796-97. While the Burgener reference noted "problems associated with the prior art designs," it did not have, as MacMillan does here, an explicit teaching away from using them. *Id.* at 800. The Der reference, into which the Burgener switch was proposed to be incorporated, also did not have any explicit teachings away from using the Burgener switch. *Id.* This is far from the case here, where MacMillan specifically teaches away from using a more complex, large, and costly processor, like a Dockser FPP, as one of its PEs.

In *Corephotonics, Ltd. v. Apple Inc.*, No. 2020-1961, 2021 WL 4944471, at *5-6 (Fed. Cir. 2021), the Court even noted as part of its holding that "[n]othing [in the references] 'criticize[s], discredit[s], or otherwise discourage[s]' investigation

into, so as to teach away from, [the proposed modification]," which is in contrast to this case. Likewise, in *In re Mouttet*, the Court found substantial evidence supporting the Board's determination that there was no teaching away and noted that its decision to invalidate a patent was required by the standard of review. 686 F.3d 1322, 1334 (Fed. Cir. 2012) ("While Mouttet's reading of Falk is plausible, our standard of review mandates that we uphold factual findings that are supported by substantial evidence as opposed to revisiting them de novo."). *See In re Kahn*, 441 F.3d 977, 990 (Fed. Cir. 2006) ("Nothing in Stanton can be said to discourage a person having ordinary skill in the art from using the visual-input control taught in Garwin in the claimed combination or to lead the skilled artisan in a direction divergent from the path taken by Kahn."). Here, unlike each case cited by Google in arguing that the Board should have given more weight to the supposed benefits of Dockser and MacMillan, one of the combination's references (MacMillan) expressly teaches away from using a design that is taught by the other of the references (a more complex Dockser FPP).

In any event, to the extent that Google argues the benefits of the combination would have outweighed the disadvantages, Google had the burden to prove that fact to the Board's satisfaction but failed to do so.

VI. THE BOARD'S FINDING THAT A POSA WOULD NOT HAVE BEEN MOTIVATED TO COMBINE A CUSTOMIZED VERSION OF DOCKSER IN THE DOCKSER-MACMILLAN-TONG COMBINATION WAS SUPPORTED BY SUBSTANTIAL EVIDENCE

The Board properly found that a POSA would not have been motivated to modify Dockser to remove its selectable and traditional-precision capabilities and then combine that modified Dockser with MacMillan and Tong to meet the claims under Ground 4 of the Petition. The Board's decision is supported by substantial evidence, including the disclosures of Dockser, Tong, and MacMillan, as well as the testimony of Dr. Khatri and Mr. Goodin.

A. Dockser Teaches a Processor that Operates at Traditional Precision with the Ability to Select a Lower Precision Based on the Application to Be Executed

As the Board properly determined, the capability to operate at *both* traditional precision or at a selected lower precision is a fundamental feature of both Dockser and Tong. Appx80; Appx164; Appx251-252.

Dockser discloses a traditional-precision floating point processor ("FPP," which is an execution unit) that allows for setting a lower precision. Dockser's teaching of an execution unit that is always capable of operating in both traditional-precision mode (its default) and reduced-precision modes was motivated by the prevalent view in the prior art that while low-precision operation might be acceptable for "certain applications," for a general-purpose processor traditional-precision

capability is "needed." Appx1470; Appx1471, ¶¶ [0003, 0018]; Appx1474-1475, claims 1, 8, 15, 20; Appx4967, ¶ 46; Appx7890, ¶ 46; Appx9668, ¶ 46.

The components of Dockser's FPP are again shown in Figure 1, reproduced below:

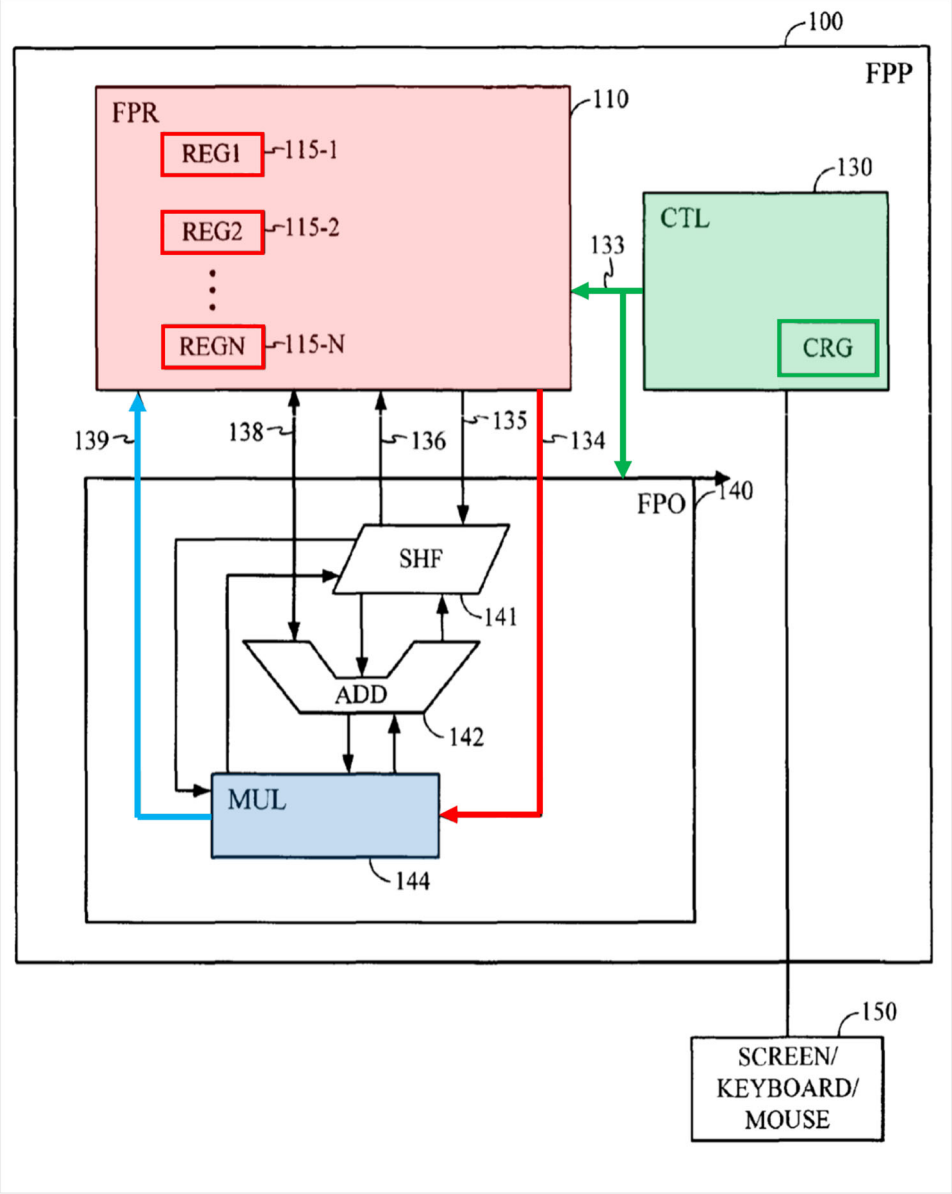

Dockser, Fig. 1 (Annotated)

Dockser's FPP has a register file (FPR 110, shaded in **red**) comprising registers (red boxes) that hold IEEE 32-bit traditional-precision values. It also includes a controller (CTL 130, shaded in **green**) with a control register (**green** box) that can store "subprecision select bits" corresponding to the desired level of precision. Appx1466, Fig. 1; *see also* Appx1471, ¶¶ [0017-18]; Appx4967-4969, ¶ 47; Appx7890-7892, ¶ 47; Appx9668-9670, ¶ 47. The FPP performs arithmetic operations (*e.g.*, including addition and multiplication) at traditional-precision – or at the desired sub-precision using values stored in the register file as operands. *See* Appx1471-1472, ¶ [0019]; Appx4967-4969, ¶ 47; Appx7890-7892, ¶ 47; Appx9668-9670, ¶ 47.

If no sub-precision is selected, Dockser's FPP performs arithmetic at "maximum precision" (*i.e.*, 32-bit IEEE traditional-precision). *See* Appx1474-1475, claims 1, 8, 15, 20; *see also* Appx1471, ¶ [0018] ("the floating-point controller 130 *may* be used to select the subprecision of the floating-point operations")); Appx4969, ¶ 48; Appx7892, ¶ 48; Appx9670, ¶ 48. Otherwise, the operation is performed at a reduced precision. *Id.* In reduced-precision modes, Dockser's FPP reduces the mantissa size by powering-down some of the cells in the register file that store the mantissa, and also powering down parts of the arithmetic logic circuits (contained, *e.g.*, within the multiplier, shaded in **blue** above) that are not needed for the selected

subprecision. Appx1472, ¶¶ [0026-27]; Appx4969, ¶ 48; Appx7892, ¶ 48; Appx9670, ¶ 48.

Nowhere does Dockser teach, or even suggest, removing its traditional-precision capabilities. Appx4997, ¶ 114; Appx7920, ¶ 114; Appx9698, ¶ 114; Appx4742, 38:3-8. A Dockser FPP is described as always needing to support a range of selectable precisions including traditional precision. Adjusting Dockser by removing its traditional-precision capacity violates a central tenet of Dockser—to always be able to execute traditional-precision operations. Appx4997, ¶ 114; Appx7920, ¶ 114; Appx9698, ¶ 114; Appx1470, ¶ [0003]. Indeed, Google's own expert, Mr. Goodin, explained that "Dockser *requires* selectable subprecision." Appx4746, 56:2-12; 57:4-12. *See also* Appx4742, 38:3-8 ("I don't believe Dockser discloses any embodiments without floating-point – without selectable precision."). Mr. Goodin also explained that he did not provide any opinions where Dockser was limited to a single subprecision. Appx4749-4750, 69:17-70:18 ("That certainly isn't part of my opinions in my declaration.").

Google's argument that "traditional-precision capability" is not fundamental to Dockser is based on a misinterpretation of a Dockser paragraph. Blue Br. at 58. Dockser states that some of its registers may be "formatted differently from IEEE 32-bit single format." Appx1471, ¶ [0017]. In the IPR, Google argued that this passage indicated that "a POSA would have been motivated to implement Dockser's

FPPs in the embedded signal-processing system with smaller than 32-bit registers." Appx721; Appx6428; Appx8405.

It turned out that Google had misinterpreted the passage. As Mr. Goodin admitted, all of Dockser's registers and multipliers are 32-bits wide. Appx4739, 26:16-23; Appx4742, 38:18-39:15; *E.g.*, Appx1471, ¶ [0017] ("Each register location 200 is configured to store a 32-bit binary floating-point number . . . ."); Appx1471-1472, ¶ [0019] (describing multiplication on the 32-bit numbers stored in the registers). Instead, the different formatting allows for 32-bit floating point numbers to be stored in *formats* (*e.g.*, a different syntax for representing of a 32-bit number, including how fractions, signs, and exponents are encoded or otherwise specified) different from that of an IEEE 32-bit format, and not, as Google was arguing, in *physically* smaller registers. Appx1471, ¶ [0017] (referring to changing from IEEE 32-bit to IEEE 64-bit double (*i.e.* using 2 32-bit registers) format).[4]

On appeal, Google changes its argument. Now it argues that the above passage of Dockser means that "the 'register file' need not be designed to be able to store numbers with 23 mantissa bits, and instead could be designed differently." Blue Br. at 58. By not previously making this argument however, Google has waived it before this Court. *In re Nuvasive, Inc.*, 842 F.3d 1376, 1380 (Fed. Cir.

[4] Dockser's statement regarding a "different number of register locations" refers to increasing or decreasing from 16 instances of registers that are 32-bits wide to, for example, 8 instances of registers that are 32 bits wide. *Id.*

2016) ("[W]e have held that a party waives an argument that it 'failed to present to the [PTAB]' because it deprives the court of 'the benefit of the [PTAB]'s informed judgment.'" (quoting *In re Watts*, 354 F.3d 1362, 1367 (Fed. Cir. 2004)). In addition, the argument that Dockser's statement relates to mantissa widths is speculative, at best.

Google also argues, for the first time, that the Board should have "accounted for" Dockser's rejection of prior-art processors that could be "designed and built to that reduced precision." Blue Br. at 59. But Google did not cite this statement to the Board, much less argue before the Board that the patent claims were obvious over these hypothetical processors or that the statement would motivate one to build a reduced-precision Dockser unit. The Board is not expected to comb through the evidence to respond to arguments that Google neglected to raise, and Google cannot now raise them for the first time on appeal. *Nuvasive*, 842 F.3d at 1380. Moreover, Google provides no evidence that a POSA would modify Dockser to implement the reduced-precision-only processor—*i.e.*, excluding traditional-precision capability—after Dockser explicitly rejected that approach.

B. Tong Does Not Support the Dockser-Tong-MacMillan Combination Because It Teaches that Full Precision Is Necessary

Similarly, Tong admits that there are "scientific programs" that "require a huge amount of precision" (Appx1482). Moreover, Tong broadly teaches that "it

appears *inevitable* that some fraction of *our operands* will require full IEEE-standard precision." Appx1483; Appx4996, ¶ 111; Appx7919, ¶ 111; Appx9697, ¶ 111. Even where Tong teaches using separate traditional-precision and reduced-precision FP units, it teaches "simply including both full and reduced precision FP units and using appropriate sleep-mode circuit techniques to shut down *the unused unit*," which amounts to pairing each low-precision FP unit with a traditional-precision FP unit in a 1:1 ratio. Appx1485; Appx4971, ¶ 54; Appx7894, ¶ 54; Appx9672, ¶ 54. Pairing each low-precision unit with a traditional-precision execution unit in a 1:1 ratio requires more physical space than a traditional-precision execution unit alone. *Id.* Indeed, this approach is presented as an option in situations where silicon real-estate is not at a premium. *See* Appx1485 ("Given the decreasing cost of silicon area . . . ."); Appx4971, ¶ 54; Appx7894, ¶ 54; Appx9672, ¶ 54.

Thus, like Dockser, Tong teaches systems that always have *both* traditional- and reduced-precision capabilities in equal measure. *See*, *e.g.*, Appx1485 (even when describing a device that has reduced precision units, it describes that system as "including *both* full and reduced precision FP units and using appropriate sleep-mode circuit techniques to shut down the unused unit."); Appx4996, ¶ 111; Appx7919, ¶ 111; Appx9697, ¶ 111. Consistent with its discussion that "our operands" will require traditional-precision, no passage in Tong discloses using only low precision modes, or even a smaller range of lower precision modes. The Board

properly recognized as much and its finding, based on the clear language of Tong, is supported by substantial evidence.

Google argues that the Court should reweigh the evidence, discard the Board's interpretation of Tong, and adopt Google's misinterpretation of Tong as suggesting a customized low-precision processor for certain applications. Blue Br. at 62-63. But Tong is not directed to such a device. Instead, Tong consistently describes a device that always includes full- and reduced-precision capabilities, or a 1:1 ratio of traditional-precision and low-precision devices, to meet the needs of a variety of applications. Appx1485.

Even if, contrary to all of the evidence presented above, the Board *could* have somehow agreed with Google's interpretation, the question here is not whether a fact-finder could have agreed with Google but whether "a reasonable mind might accept the evidence to support the finding." *Redline Detection,* 811 F.3d at 449. As shown above, the Board's reasoning is well-supported by the evidence and, therefore, should be affirmed.

C. A POSA Would Not Modify Dockser as Google Suggests

Google now argues that a POSA would ignore the disclosures of Tong and Dockser and be motivated to modify Dockser to remove "unnecessary circuitry." Blue Br. at 64-65. But Google's brief does not identify any evidence on the record (other than Mr. Goodin's speculation—*see* Appx1302, ¶ 394; Appx7156, ¶ 396;

Appx8926, ¶ 423) identifying any circuitry that would be removed. *See generally* Blue Br.

To the extent that Google returns to its argument that Dockser's discussion of storing numbers in formats other than IEEE single precision suggests that Dockser endorses specialized physically smaller-than-32-bit registers, that argument is foreclosed by Dockser's command otherwise. Appx1471, [¶ 0017] ("Each register location 200 is configured to store a 32-bit binary floating-point number . . . ."); Appx1471-1472, ¶ [0019] (describing multiplication on the 32-bit numbers stored in the registers).

Moreover, the Board properly found that MacMillan teaches away from a Dockser-Tong-MacMillan device because such a device would add complexity through its "customized implementation of the multiplier logic in Dockser's FPP." Appx82; Appx166-167; Appx254. Thus, even if there were some motivation to modify Dockser to remove its traditional-precision capabilities despite its own teachings to the contrary, the precision selection circuitry would nevertheless add complexity and conflict with MacMillan's express teaching to reduce cost by reducing complexity, as discussed above. *Id.*

Substantial evidence in the record supports this finding. Dr. Khatri explained in the IPR proceedings that creating a customized version of Dockser would increase complexity. For example, even if it were possible to provide a version of Dockser

that only selected between a limited set of low-precision settings, as Google now argues, Dr Khatri testified that the Dockser FPP would still be more complex than a conventional execution unit due to the need for the selection circuitry. Appx4998, ¶ 115; Appx7921, ¶ 115; Appx9698-9699, ¶ 115. Moreover, Dr Khatri testified that Google's combination would require specialized customized registers, logic elements, arithmetic units, and programming models. *Id.* Given these additional costs and increased complexity, a POSA would not be motivated, except by hindsight, to remove Dockser's traditional-precision capability and create a crippled Dockser limited to some still-unspecified range of precisions, contrary to the teachings of Dockser itself requiring continual support for traditional-precision numbers and arithmetic. Appx4996-4998, ¶¶ 110-15; Appx7919-7921, ¶¶ 110-15; Appx9696-9699, ¶¶ 110-15.

In the face of this reasoning, Google nonetheless argues that a POSA would be motivated to remove Dockser's traditional-precision capability solely to achieve some sort of power or space savings. Blue Br. at 56. In particular, Google states that "POSAs would have customized Dockser's FPPs by removing register elements that would have stored mantissa bits that the processor would never use, because having register elements that would never be used would 'waste circuit space or incur unnecessary cost.'" *Id.* (citing Appx1302; Appx7156; Appx8926.)

Google did not raise this argument in the IPR, and for good reason. Most critically, there is no evidence in the record (other than Mr. Goodin's unsupported speculation—*see* Appx1302, ¶ 394; Appx7156, ¶ 396; Appx8926, ¶ 423) supporting Google's new argument. The cited portion of Mr. Goodin's declaration states only that Dockser's register file may be "***formatted differently***" and that, according to Mr. Goodin, means that those registers may be smaller than 32-bits. Appx1302, ¶ 394; Appx7156, ¶ 396; Appx8926, ¶ 423. As discussed above, and as Google tacitly admits elsewhere, the cited portion of Dockser however refers only to ***formats*** and does not suggest or imply that Dockser would function with smaller than 32-bit registers. *See supra*, § II.B (citing Appx4739, 26:16-23; Appx4742, 38:18-39:15); Appx1471, ¶ [0017] ("Each register location 200 is configured to store a 32-bit binary floating-point number"); Appx1471-1472, ¶ [0019] (describing multiplication on the 32-bit numbers stored in the registers)).

Further, a POSA would not design a MacMillan device with a hundred or more Dockser FPPs, regardless of whether they were traditional-precision Dockser FPPs, or Dockser FPPs "modified" by physically hacking out register elements from traditional-precision Dockser FPPs. As the Board recognized, and as explained above, any Dockser FPP is more complex than a conventional execution unit. Thus, placing hundreds of Dockser FPPs (even those with artificially hacked out register elements) would still increase complexity with every addition of such a more

complex Dockser variable precision FPP, in contravention of MacMillan's express direction to minimize costs by minimizing complexity. It would not be obvious to make such a modification.

Accordingly, a POSA would not make Google's proposed Dockser/MacMillan/Tong combination.

D. Google Relied on Economic Efficiencies as the Motivation to Combine While Simultaneously Arguing that Economic Concerns Are Irrelevant as a Matter of Law

Google argues that the Board's statement that the customized Dockser FPP would add complexity justifies reversal. Blue Br. at 66. Google then cites to its expert's declaration stating that customization typically reduces the size of a component. *Id.* at 66-67.

Specifically, Google argues that the "customized" version would save costs by reducing extra circuits and waste. However, Google does not dispute that the customized version would still require the additional control circuits for enabling variable precision, as the Board noted. Additional control circuits for enabling variable precision add complexity. Appx74; Appx158; Appx244-245. Thus, the Board's decision is still supported by substantial evidence. Despite having the burden of persuasion, Google did not provide evidence showing that its hypothetical modifications to Dockser would result in any overall cost savings.

Moreover, Google's arguments on this issue are predicated on the use of embedded systems specifically designed to execute applications that require only low-precision processing, as opposed to generally programmable systems, such as personal computers, to realize some amount of cost savings through power savings. Blue Br. at 66-67. But in its written submissions to the Board, Google's motivation to combine was not based on embedded systems, but on the idea that a POSA would seek to save power in a laptop personal computer implementation. *Id.* at 39 ("MacMillan's embodiments include battery-powered devices like 'laptops' and Dockser teaches that power consumption is 'of particular concern in battery operated devices' like 'laptops.'"). The result is a Frankenstein-like combination that combines features from completely different sorts of systems—embedded (non-programmable) systems focused on efficiency and laptop personal (programmable) computers focused on minimizing power consumption— and then attempts to combine that combination with MacMillan's supercomputer focused on low cost and simplicity. Google did not argue and presented no evidence to the Board in support of such a confusing combination.

VII. THE BOARD'S DECISION IS FURTHER SUPPORTED BY STRONG EVIDENCE OF OBJECTIVE INDICIA OF NON-OBVIOUSNESS

The Board reached the issue of objective indicia only for claims covering a single LPHDR execution unit (the "at-least-one" LPHDR execution unit claims),

which it referred to as the "Dockser-/Tong-challenged Claims." (Appx36; Appx124; Appx209).). The Board did not reach this issue with respect to the claims at issue because its findings regarding motivation to combine was dispositive for those claims: "we need not assess objective indicia of nonobviousness with respect to [the "greater-than one" claims]." Appx36 n.4; Appx124 n.3; Appx209 n.3.

Substantial evidence supports the Board's finding of no motivation to combine, without considering objective indicia. If this Court were to disagree with the Board on that point, as well as the alternate basis for patentability presented in Section IV, the appropriate course would be to remand for consideration of the evidence of objective indicia regarding the claims at issue, because such evidence of, when offered, must always be considered. *Apple Inc. v. Samsung Elecs. Co.*, 839 F.3d 1034, 1048 (Fed. Cir. 2016) (*en banc*) ("Objective indicia of nonobviousness ***must*** be considered in every case where present.").

Although the Board concluded Singular's evidence of objective indicia did not save the "at-least-one" LPHDR execution unit claims, most of its reasons for that finding do not apply to the claims at issue, all of which involve multiple LPHDR execution units.

There was substantial evidence of the following objective indicia of non-obviousness for the claims at issue:

- Google's own initial skepticism of the claimed inventions, expressed by its in-house experts;
- Praise from leading experts in the field, including Google's own employees; and
- Google's copying of the invention.

A. Google's Initial Skepticism of the Invention Shows Non-Obviousness

"Evidence of industry skepticism weighs in favor of non-obviousness." *WBIP, LLC v. Kohler Co.*, 829 F.3d 1317, 1335-36 (Fed. Cir. 2016); *see also e.g.*, *Neptune Generics, LLC v. Eli Lilly & Co.*, 921 F.3d 1372, 1377-78 (Fed. Cir. 2019) ("Evidence of industry skepticism is a question of fact that weighs in favor of non-obviousness.").

The skepticism of Google's engineers after the inventor, Dr. Bates, disclosed the invention to Google in 2010 supports non-obviousness. Appx4671; Appx4049. Dr. Bates taught Google that a computer system covered by the claims at issue enables massive scale and parallelism in the form of a large number of LPHDR execution units, compared to IEEE traditional 32 bit floating-point execution units. Appx4052; Appx4053; Appx4054-4059; *see also* Appx263-264.

In response to Dr. Bates's disclosures, Google's engineers expressed skepticism. Appx4682; Appx4161; Appx4673; Appx4676; Appx4678.

Google points out that with respect to the "at-least-one" LPHDR execution unit claims, the Board found Singular's skepticism evidence did not "relate to whether low-precision operations 'would be commercially valuable'" and did "not show skepticism as to whether or not or how a problem could be solved or whether such applications would *work*." (emphasis in original). Blue Br. at 69. The Board, however, analyzed Singular's skepticism evidence only through the prism of the "at-least-one" LPHDR execution unit claims. When the Board indicated that Singular's skepticism evidence was directed more towards the invention's supported applications than towards whether the invention would work in the first place, the Board was focused on the "approximate nature" of the operations (*i.e.*, pertaining to the claim elements of the at-least-one LPHDR execution unit claims, and not the multiple-LPHDR execution unit claims at issue). Appx39; Appx128; Appx213-214. Focusing on the claims at issue, the evidence shows Google to have been skeptical about whether practicing Singular's technology would even work. Appx4682.

The skepticism of Google's engineers is evidence that the claims at issue would not have been obvious because there is a nexus between these statements and the claims at issue. Appx5001, ¶ 123; Appx7924, ¶ 123; Appx9702-9703, ¶ 123.

As explained above, Google's expressed skepticism about whether a computer featuring a traditional-precision execution unit and a much larger number

of LPHDR execution units configured would ***work correctly***, not just about the range of applications that would benefit from Dr. Bates's invention.

B. Praise from Leaders in the Field Shows Non-Obviousness

Praise of the invention by others is another objective indicator of non-obviousness. *Mintz v. Dietz & Watson, Inc.*, 679 F.3d 1372, 1379 (Fed. Cir. 2012).

As Dr. Bates's invention circulated more broadly within Google in 2011, the initial skeptical reaction to that invention gave way to praise from some of the leading figures in computer science. Appx4949; Appx4687; Appx4688; Appx4689. In response to a detailed description from Dr. Bates, **[name]**, a world-renowned scientist in artificial intelligence, prophetically called it "**[praise statement]**" and "**[praise statement]**." Appx4689 (emphasis in original).

Google argues this evidence does not indicate any specific ideas or features that were allegedly praised. Blue Br. at 70. The Board's analysis was, again, viewed through the prism of the "at-least one" claims. With respect to the claims at issue, however, there is a clear nexus between the invention as described to Mr. **[name]** and Mr. **[name]** praise. Appx4689. Mr. **[name]**, a renowned artificial intelligence expert, was referring to this entire computer system when he said, "**[praise statement]**."

C. Google's Copying of the Invention Shows Non-Obviousness

"That a competitor copied technology suggests that it would not have been obvious." *WBIP*, 829 F.3d at 1336.

As described above, Dr. Bates met with Google and disclosed his invention, as covered by the claims at issue. Appx4690; Appx4051-4059; Appx5001, ¶ 123; Appx7924, ¶ 123; Appx9702-9703, ¶ 123; Appx4081; Appx4094; Appx4113.

As Dr. Bates was describing his invention to Google, Google had identified a "daunting and scary" problem: to use artificial intelligence for speech recognition, it would need to "double the computing footprint of Google just to support, like, a slightly better speech recognition model for a modest fraction of [its] users." Appx4696. Google solved this problem by copying Dr. Bates's invention. After receiving detailed descriptions of the invention (Appx4049; Appx4061; Appx4081; Appx4094; Appx4113; Appx4689), the invention, as both described by Dr. Bates and covered in the claims at issue, was incorporated into the TPUv2 and TPUv3 products. Appx5004-5014, ¶¶ 128-153; Appx7927-7937, ¶¶ 128-153; Appx9704-9715, ¶¶ 128-153.

Dr. Khatri performed a detailed analysis showing that the TPUv2/3 incorporates each and every element of the claims at issue. Appx9704-9715, ¶¶ 128-153. *See WBIP*, 829 F.3d at 1329 ("[T]here is a presumption of nexus for objective considerations when the patentee shows that the asserted objective evidence is tied

to a specific product and that product 'is the invention disclosed and claimed in the patent.'")

In addition to this presumption, more strong evidence shows Google's copying. Appx4113; Appx4718; Appx4131; Appx4286; Appx4730; Appx4706; Appx4709, ¶ 8; Appx4713; Appx4716.

Taken together:

(1) Google's receipt of descriptions of the invention, as set out in detail in Appx4049; Appx4061; Appx4081; Appx4094; Appx4113; Appx4689 and Appx4182; Appx5001, ¶ 123; Appx7924, ¶ 123; Appx9702, ¶ 123,

(2) Google's incorporation of the invention in its TPUv2/3 products, as set out in detail in Dr Khatri's testimony, and

(3) and the remainder of the evidence discussed above,

collectively amount to substantial evidence of copying by Google.

In its discussion of "copying and commercial success" with regard to the "at-least-one" claims, the Board reached the following findings specific to copying: 1) A nexus cannot be presumed because the Google product had other significant features that are not part of the claims which were viewed as critical, successful features that were reasons for the products' successful performance; 2) Low-precision arithmetic was not novel and thus cannot be the basis for a presumption of nexus; 3) Singular had not identified and explained a specific combination of elements that could serve

as a nexus for the evidence of copying; and 4) Singular "does not even attempt to demonstrate that [Google's products] "perform the same bit-dropping techniques as Dockser." Appx50; Appx138; Appx224.

As to the last Board finding, Google misquotes the Board as having stated that Singular did not "even attempt to demonstrate" that Google's [TPU] products *met the challenged claims' 'X and Y percentage [error] limitations.*'" Blue Br. at 69. The Board actually was discussing whether or not the Google products "perform the same bit-dropping techniques as Dockser." Appx50; Appx138; Appx224. This distinction is important because Singular did not need to demonstrate that TPUv2/3s perform the same bit dropping techniques as Dockser to prove that Google's TPUv2/3s meet the X/Y limitations, as will be explained on remand, if that occurs. Part of that explanation will Google's description of the TPUV2/3s, included as part of Singular's analysis of the X/Y claim limitations: "While its inputs and outputs are 32-bit floating point values, the MXU typically performs multiplications at the reduced precision of bfloat16. . . . Multiplication performed at bfloat16 precision uses 7 bits for the mantissa." Appx5008, ¶ 135; Appx7931, ¶ 135; Appx9708-9709, ¶ 135. Google's admission rendered any need to refer to "bit dropping techniques" as moot, since it is apparent from Appx4939 that TPU v2 and v3 computers reduce the bit count of their inputs from 32 bits to 16 bits (*i.e.*, akin to "bit dropping techniques") at the point of multiplication.

As to the first three Board findings specific to copying, none apply to the claims at issue because of differing claim scope. Accordingly, the Board would need to consider that nexus in the first instance on remand.

As with all the other limitations of the claims at issue, Singular presented substantial evidence of copying by Google in making the TPUv2/3 computer systems.

Google's copying evidences that the claims at issue would not have been obvious. *E.g.*, *Liqwd, Inc. v. L'Oreal USA, Inc.*, 941 F.3d 1133, 1139 (Fed. Cir. 2019) (disclosure of patented method and later adoption of technology supported finding of copying); *see also Power Integrations, Inc. v. Fairchild Semiconductor Int'l, Inc.* 711 F.3d 1348, 1369 (Fed. Cir. 2013).

VIII. ALTERNATIVELY, THE COURT SHOULD AFFIRM THE BOARD'S DECISION AS TO GROUNDS 3 AND 4 BECAUSE THE COMBINATION DOES NOT DISCLOSE OR RENDER OBVIOUS THE "EXCEEDS" LIMITATION

If the Court determines that the Board erred in finding no motivation to combine Google's references, the decision as to Grounds 3 and 4 should, nevertheless, be affirmed as to most of the claims at issue because Google cannot show the combination discloses or renders obvious claims 3 and 23 of the '961 Patent, claims 3-8 of the '156 Patent, and claims 5-6, 8, 10-18, 35-61, 63, and 67-70 of the '273 Patent ("the Exceeds Claims").

Each of these claims includes a limitation that the number of LPHDR execution units exceeds the number of "execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide." (Each of the latter units are referred to as a "32-bit traditional-precision execution unit," and the claim limitation as the "exceeds limitation"). *E.g.*, Appx288 at cl. 3; Appx4985, ¶ 85; Appx7908, ¶ 85; Appx9686, ¶ 85.

The Board did not reach this issue because its finding that there was no motivation to combine Dockser and MacMillan (or Dockser and Tong and MacMillan) was dispositive. Appx77 n.8; Appx161 n.8; Appx247 n.8. However, the Board's decision of non-obviousness as to the Exceeds Claims would be equally supported under this alternative basis.[5]

[5] This alternative basis is relevant to the Board's decision as to Ground 3 and Google's primary argument under Ground 4. Google's "alternative" argument under Ground 4 (*i.e.*, a Dockser/MacMillan/Tong device where Dockser's full-precision capability has been removed) is deficient, as described in Section II, *supra*.

A. The Dockser-MacMillan Combination Does Not Disclose Systems Where the Number of LPHDR Execution Units Exceeds the Number of Execution Units Adapted to Perform Full-Precision Multiplication

The Dockser-MacMillan combination consists of an array of PEs where the MacMillan PEs are replaced by Dockser FPPs, *i.e.*, FPPs that can perform an operation using *either* 32-bit precision (*i.e.*, traditional precision) or some reduced precision using fewer than 32 bits. The Board construed "LPHDR execution unit" in all the claims as broad enough to include a Dockser FPP, because Dockser's execution unit could be flexibly adapted to perform an LPHDR operation, notwithstanding it could alternatively perform a 32-bit traditional-precision operation. Appx17; Appx105; Appx191.

As shown below, a Dockser FPP is adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide and is thus a 32-bit traditional-precision execution unit. *Id.* Therefore, assuming Dockser's FPP to also be an LPHDR execution unit, as the Board ruled, the Dockser/MacMillan combination would not meet the claim's "exceeds limitation" because that same Dockser LPHDR execution unit would ***also*** be a 32-bit traditional-precision execution unit (*i.e.*, "adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide"). Each addition of a Dockser FPP to a MacMillan computer would amount to the addition of both an LPHDR execution unit and a 32-bit traditional-precision execution unit.

Thus, the number of LPHDR execution units in the Dockser/MacMillan combination can never exceed the number of 32-bit traditional-precision execution units. Appx4985-4986, ¶ 86; Appx7908-7909, ¶ 86; Appx9686-9687, ¶ 86.

1. Dockser's FPP Is a Traditional-Precision Execution Unit Because It Is Adapted to Execute at Least the Operation of Multiplication on Floating Point Numbers that Are at Least 32 Bits Wide

The FPP disclosed in Dockser is an execution unit. Appx676; Appx6383; Appx8358. And Google does not dispute the Dockser FPP is adapted to execute the operation of multiplication on floating point numbers that are at least 32 bits wide under the plain meaning of the term.[6] Appx5054-5057; Appx7976-7979; Appx9755-9758.

Dockser discloses that its FPP takes input operands that are 32 bits wide and stores them in IEEE-754 floating point 32 format. *See e.g.*, Appx1471, ¶ [0017]; Appx4988, ¶ 92; Appx7911, ¶ 92; Appx9689, ¶ 92. Dockser also discloses that its FPP includes "a floating-point multiplier (MUL) 144 configured to execute floating-point multiply instructions." Appx1471-1472, [¶ 0019]; Appx4988, ¶ 92; Appx7911, ¶ 92; Appx9689, ¶ 92. These points are not in dispute. Appx5044; Appx7966; Appx9745; Appx4741 at 36:16-21; Appx4742 at 38:18-39:15.

[6] Both parties agreed that "adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide" means that the execution unit performs multiplication in "traditional"—*i.e.*, full-precision. Appx717; Appx6442; Appx8395.

Nothing more is required for a Dockser FPP to be "adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide." *See In re Man Mach. Interface Techs. LLC*, 822 F.3d 1282, 1286 (Fed. Cir. 2016) ("[T]he phrase 'adapted to' generally means 'made to,' 'designed to,' or 'configured to,' though it can also be used more broadly to mean 'capable of' or 'suitable for.'"); Appx4988, ¶ 92; Appx7911, ¶ 92; Appx9689, ¶ 92.

Therefore, the Dockser/MacMillan device cannot satisfy the "exceeds" limitation. Appx4988-4989, ¶ 93; Appx7911-7912, ¶ 93; Appx9689-9690, ¶ 93. In one example, the Dockser/MacMillan device would have 256 LPHDR execution units (counting each Dockser FPP as an "LPHDR execution unit"). However, because each Dockser FPP is also "adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide," the Dockser/MacMillan device would also have *at least* 256 execution units that satisfy that criterion. *Id.* As a result, the Dockser/MacMillan device would not meet the "exceeds limitation." *Id.*

2. Google's Interpretation of "Adapted to Execute at Least the Operation of Multiplication on Floating Point Numbers that Are at Least 32 bits Wide" Is Inconsistent with the Ordinary Meaning of the Exceeds Limitation and the Specification

A construction of "adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide" that would exclude the Dockser FPPs would be inconsistent with the ordinary meaning of the claim. Appx4989, ¶¶ 94-95; Appx7912, ¶¶ 94-95; Appx9690, ¶¶ 94-95. A POSA would understand this term to include all units in the device that are designed to perform (or are capable of performing) "multiplication on floating point numbers that are at least 32 bits wide," even if those units can also perform other operations at lower precision. Appx4989, ¶ 95; Appx7912, ¶ 95; Appx9690, ¶ 95.

Google argued to the Board that a component found to be an "LPHDR execution unit" (*i.e.*, one element of the claim), cannot also be a traditional-precision multiplication execution unit (, another element of the claim). Appx717; Appx6442; Appx8395. That is, that one device—the Dockser FPP—cannot meet two limitations in a claim. But the same structure *can* meet two limitations. *See, e.g., Applied Med. Res. Corp. v. U.S. Surgical Corp.*, 448 F.3d 1324, 1333 n.3 (Fed. Cir. 2006) ("[T]he use of two terms in a claim requires that they connote different *meanings,* not that they necessarily refer to two different *structures.*") (emphasis in original). Therefore, the Dockser FPP can be both an "LPHDR execution unit" and

a traditional-precision multiplication execution unit, for the reasons explained above.

Google's argument imports an unsupported negative limitation into the claim; that is, Google argues that "execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide" actually means "execution units in the device *that are not LPHDR execution units* and that are adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide." There is no support for this negative limitation, anywhere in the intrinsic record.

In sum, under the plain language of the Exceeds Claims, a Dockser FPP is both an LPHDR execution unit and an "execution unit[] adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide." Accordingly, the number of Dockser FPPs cannot exceed the number of "execution units adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide." The combination thus cannot render the Exceeds Claims obvious.

CONCLUSION

This Court should affirm the Board's decision finding the claims at issue not unpatentable as obvious.

Dated: December 23, 2022

Respectfully submitted,

/s/ Alfred R. Fabricant

Alfred R. Fabricant
Peter Lambrianakos
Vincent J. Rubino, III
Enrique W. Iturralde
FABRICANT LLP
411 Theodore Fremd Avenue
Suite 206 South
Rye, New York 10580
Telephone: (212) 257-5797
Facsimile: (212) 257-5796

James J. Foster
PRINCE LOBEL TYE LLP
One International Place,
Suite 3700
Boston, MA 02110
(617) 456-8000

Attorneys for Appellee
Singular Computing LLC

CERTIFICATE OF COMPLIANCE

I certify that this Brief complies with the type-volume limitation of Fed. R. App. P. 27(d)(2)(A). This Brief contains 13,310 words, excluding the parts of the Brief exempted by Fed. R. App. 27(d) and Fed. Cir. R. 27(d). This Brief complies with the typeface requirements of Fed. R. App. P. 32(a)(5) and the type style requirements of Fed. R. App. 32(a)(6). This Brief has been prepared in a proportionally spaced typeface using Microsoft Word for Office 365 in 14-point Times New Roman font.

Dated: December 23, 2022

Respectfully submitted,

/s/ Alfred R. Fabricant

Alfred R. Fabricant
Peter Lambrianakos
Vincent J. Rubino, III
Enrique W. Iturralde
FABRICANT LLP
411 Theodore Fremd Avenue
Suite 206 South
Rye, New York 10580
Telephone: (212) 257-5797
Facsimile: (212) 257-5796

James J. Foster
PRINCE LOBEL TYE LLP
One International Place
Suite 3700
Boston, MA 02110
(617) 456-8000

Attorneys for Appellee
Singular Computing LLC

CERTIFICATE OF SERVICE

I hereby certify that I caused to be electronically filed the foregoing Corrected Non-Confidential Response brief of Appellee, Singular Computing LLC with the Clerk of the Court for the United States Court of Appeals for the Federal Circuit by using the appellate CM/ECF system on February 3, 2023 and thus caused to be served on all registered counsel of record a copy of the same via the CM/ECF system.

Dated: February 3, 2023

/s/ Alfred R. Fabricant
Alfred R. Fabricant